LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-470873

2010 Computation Directorate Annual Report

D. L. Crawford

February 16, 2011

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

2010 COMPUTATION ANNUAL REPORT

LAWRENCE LIVERMORE NATIONAL LABORATORY

"The science of today is the technology of tomorrow."

Edward Teller, Director Emeritus

CONTENTS

Message from the Associate Director

DONA CRAWFORD, ASSOCIATE DIRECTOR

Thanks to founders Edward Teller, Ernest Lawrence, and the scientists and engineers they enlisted to open Lawrence Livermore National Laboratory (LLNL) in September 1952, computing is in LLNL's DNA and has been integral to Laboratory research and development (R&D), science and technology, and operations ever since. "Even before the Livermore Laboratory began, I was convinced that computers were instrumental to success. The problems we were trying to solve were too lengthy and complicated to be tackled by hand," Teller wrote in his *Memoirs: A 20th Century Journey in Science and Politics*. "We needed to have the best available computers and to develop new computer techniques." Since the Laboratory's earliest days, the driver for the development of computing capabilities has been the needs of the programs.

Sid Fernbach, an early computing pioneer at LLNL, "was the man responsible for one of the most important, if least visible Laboratory programs—computation," Teller wrote. Fernbach pursued the best computing systems available at the time with "bulldog-like tenacity," according to Teller. As a result, "it is not an overstatement to say that [Fernbach's] vision guided the commercial development of computers during the next two decades."

We have come a long way since the days of the Univac, the Laboratory's first computer system, when nothing could be more exciting than working on a machine with a 1,000-word memory. Today, we continue to build on that legacy of pioneering scientific computing. The Laboratory's Computation Directorate influences what is now a global computing enterprise by demonstrating how high-performance-computing-enabled modeling, simulation, and analysis can drive scientific and technological innovation. We accomplish this through a multi-faceted strategy that involves working with LLNL programs and directorates, other national laboratories, academia, and industry.

To maintain U.S. preeminence in high performance computing (HPC) and its application to the great scientific challenges of our time requires that we continue to push the boundaries of scientific and technological progress across a number of domains involving computer hardware, software, models, algorithms, technologies, and applications. Today we are laying the foundation for Sequoia, a 20-quadrillion-floating-point-operations-per-second BlueGene/Q system scheduled for delivery in early 2012, and we are using Dawn, a 500-trillion-floating-point-operations-per-second BlueGene/P system, to prepare applications for next-generation massively parallel computing. Since its arrival in 2009, Dawn has demonstrated groundbreaking science, better code performance, and some of the highest resolution simulations run in a variety of scientific fields, including climate modeling, laser–plasma interactions, and hydrodynamics. The development of these systems and their application is made possible by a long-term partnership with IBM.

In collaboration with other laboratories, Computation scientists are performing R&D that will enable the codesign and use of next-generation supercomputers as they are developed as part of the Department of Energy's (DOE's) Exascale Initiative. Nearer term, the National Nuclear Security Administration's (NNSA's) Advanced Simulation and Computing (ASC) program conducts applied research that provides new technology for ASC Integrated Codes and Laboratory Strategic Mission Support. These efforts allow us to turn research discoveries, supported by Office of Science and Laboratory Directed Research and Development investments, into usable products that provide immediate benefit to multiple Laboratory programs.

Our industrial partnerships are strong and are yielding results that are benefitting the entire computing community. More recently, these collaborations have started to bring high-fidelity predictive simulation capabilities to private industry where the norm for high-end computing has often been to use clusters with 128 processors or less. Computation is working with industry to bridge the capability gap by providing predictive simulation expertise and access to leading-edge computers.

In 2010, the Laboratory began to take collaboration with industry and academia to a new level with the start of the Livermore Valley Open Campus. The open campus, an area outside the perimeter fence on the Laboratory's southeast sector, will facilitate broadened LLNL collaboration with industry and academia in a wide range of fields, including computation, to the benefit of LLNL research and industrial innovation. Industry–national laboratory synergy can only accelerate the technological development needed to address national challenges and enhance the nation's economic competitiveness.

DOE is leveraging its investment in basic science to advance solutions to the nation's energy challenges, and Computation is uniquely positioned to meet the growing modeling and simulation demands of this new research. We provide modeling and simulation capabilities for one of DOE Secretary Steven Chu's three recently announced Energy Innovation Hubs, which focuses on devising ways to construct or retrofit buildings to be more energy efficient. This year, we also received funding to contribute to the DOE Carbon Capture Simulation Initiative. "This partnership will not only help mitigate climate change, it will create new jobs and position the United States as a leader in carbon capture and storage technologies […] for years to come," said Chu. To address serious challenges such as national security, environmental security, and economic security, we must continue to grow and influence modeling and simulation efforts across the complex.

In the pages of this report, you will see the myriad ways Computation is forging new frontiers, serving our customers today, and anticipating their needs for tomorrow. We push the state of the art in the development of computational tools and their application to the scientific problems at the heart of LLNL missions and national need. Our goals are straightforward: we want nothing less than to change the world and to leave no doubt about LLNL's leadership role in making that change possible.

DONA CRAWFORD
ASSOCIATE DIRECTOR, COMPUTATION

RETURN TO TOC

AN AWARD-WINNING ORGANIZATION

The stories in this annual report present a crosssection of Computation's accomplishments in research, high performance computing, software applications, and information technology and security. In addition to the projects highlighted in the report, several Computation personnel and projects received prestigious external recognition in 2010. Some of the honors are featured in this section.

Greg Bronevetsky

Early Career Research Grant

Greg Bronevetsky is dedicating his early scientific career to ensuring that the increasing power, size, and complexity of the supercomputers critical to national security research do not come at the expense of their reliability. Bronevetsky's sponsors recognized the importance of his work with a U.S. Department of Energy (DOE) Early Career Award. He will receive $500,000 a year for the next five years to conduct his research.

The focus of Bronevetsky's project is to understand the effects of hardware faults and failures on computer applications by developing an automated methodology that is more efficient and less costly than current methods. His goal is to make large-scale computing systems more usable, manageable, and reliable despite their growing complexity. Bronevetsky will develop automated methods to analyze the behaviors of hardware and software systems, which will enable a new generation of systems that can deliver high performance and productivity even on increasingly complex and unreliable supercomputing systems.

Regional Technology Award for Virus Detection Tool

Computation researchers Shea Gardner, Kevin McLoughlin, and Tom Slezak are members of a Lawrence Livermore National Laboratory (LLNL) team that captured an outstanding technology development award in the Federal Laboratory Consortium's (FLC's) Far West Region competition. Their Lawrence Livermore Microbial Detection Array (LLMDA) technology can detect more than 2,000 viruses and 900 bacteria in 24 hours. LLMDA has shown value for applications in detecting bioterrorism events, product safety, and diagnostics (see page 30).

LLNL has won 12 FLC Far West Regional Awards since 2007. FLC is the nationwide network of federal laboratories that provides a forum to develop strategies and opportunities for linking the laboratory mission technologies and expertise with the marketplace. There are more than 100 federal laboratories and facilities in the FLC Far West Region.

Outstanding Leadership in HPC

Computation Associate Director Dona Crawford was featured in *insideHPC* as a "Rock Star of HPC" in December. She was chosen for being "a true leader who inspires and motivates with vision and passion." Crawford is the first woman to receive the honor.

First Place in Cyber Security Showdown

The Cyber Security Program Network Security Team won first place in the national laboratory Cyber Tracer competition held during the Tracer FIRE 2 Workshop. Eleven cross-laboratory teams competed for two days in malware reverse engineering, network forensics, host forensics, cryptography, social engineering, and hidden data discovery.

"LEEDing" the Way in Energy and Environmental Design

The Computation Directorate is committed to continuously reducing energy costs and finding operational efficiencies in its computing facilities. The Terascale Simulation Facility (TSF) received a Leadership in Energy and Environmental Design (LEED) gold certification under the U.S. Green Building Council rating system (see page 26). LEED is an internationally recognized green-building certification system. It provides third-party verification that a building or community was designed and built using strategies aimed at improving performance in energy savings, water efficiency, carbon dioxide emissions reduction, and other factors.

TSF is a 253,000-square-foot building that houses some of the world's fastest supercomputers, including Dawn (BlueGene/P) and BlueGene/L—Advanced Simulation and

Livermore's Cyber Tracer team included John Townsend, Adam Sealey, Willem Verschuur, Erwin Lopez, Matthew Myrick, and Steve McManus.

Tina Eliassi-Rad

Gary Laguna

Jeffrey Westcott

Computing program systems largely dedicated to stockpile stewardship. TSF represents an innovative design that emphasizes function over form.

Making a Difference in Students' Lives

Tina Eliassi-Rad, Gary Laguna, and Jeffrey Westcott were named DOE Office of Science Outstanding Mentors. They were nominated by LLNL students who participated in internships during the summer of 2009.

In 2010, Computation's Institute for Scientific Computing Research hosted 105 students in the scholar program—the largest number of participants since 2004. The program included 56 Ph.D. candidates, 25 graduate students, and 24 undergraduates (and two faculty) from 60 universities across the globe. Computation is resolute in its mission to create a student internship experience that is worthwhile and rewarding. The goal is for participants to leave with an understanding of the exciting career opportunities in Computation and of the breadth, depth, and impact of Laboratory science.

The Institute for Scientific Computing Research hosted 105 students in the 2010 scholar program.

RETURN TO TOC

Research Provides a Solid Foundation for Fulfilling Laboratory Missions

DAVID L. BROWN
BROWN137@LLNL.GOV

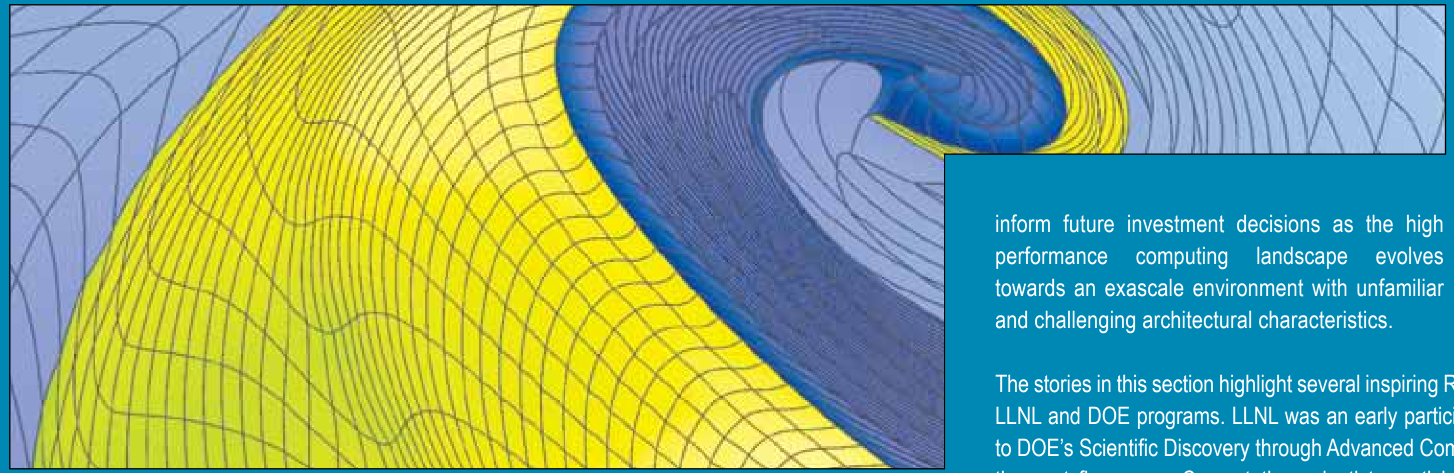

Research and development (R&D) in the Computation Directorate continually advances the state of the discipline in support of Lawrence Livermore National Laboratory (LLNL) and Department of Energy (DOE) missions. During the previous decade, Computation expanded its R&D efforts from a focus on algorithmic and software support for the Laboratory's stockpile stewardship mission to a broad R&D portfolio that also supports challenges in Global Security, the National Ignition Facility (NIF), the DOE Office of Science, and the Department of Defense (DoD). Computation's research is funded by a combination of internal support via the Laboratory Directed Research and Development (LDRD) program and external support via the National Nuclear Security Administration's (NNSA's) Advanced Simulation and Computing program, the DOE Office of Advanced Scientific Computing Research, and DoD.

Computation's R&D program encompasses activities in several disciplines, including applied mathematics, computer science, and data science. The directorate's research pursuits are guided by the Laboratory's strategic Roadmap to the Future and by external exascale computing initiatives, sponsored jointly by NNSA and the Office of Science. During the past year, Computation continued to provide expertise and leadership for the DOE Exascale Initiative workshops and reports, which will inform future investment decisions as the high performance computing landscape evolves towards an exascale environment with unfamiliar and challenging architectural characteristics.

The stories in this section highlight several inspiring R&D successes that directly benefit LLNL and DOE programs. LLNL was an early participant in and significant contributor to DOE's Scientific Discovery through Advanced Computing (SciDAC) program. During the past five years, Computation scientists participated in 17 SciDAC projects and led two of the nine SciDAC Centers for Enabling Technology. Award-winning LLNL computational technologies, developed by Computation, contributed to the scientific success of many SciDAC projects. In response to requests from intelligence community sponsors, Computation scientists developed new methods for large-scale text analysis, novel techniques for rapid processing of streaming three-dimensional video surveillance data, and they are researching eigensolvers for extremely large matrices associated with large-scale graph analysis. Computation scientists continue to improve methods for modeling and simulation in support of the stockpile stewardship program; one of the directorate's contributions is a class of new high-order, invariant-preserving finite element methods that can enhance the predictability of hydronamics simulations. Computation's compiler research effort spun off novel uses of automated code analysis, including predicting code behavior on future computer systems, optimizing code performance, and automatically converting Message Passing Interface (MPI) programs to run on multi-core architectures. Computation scientists also developed a data science approach for the discovery and prediction of damage in optics elements for NIF, a contribution that minimizes the downtime of the facility by anticipating optics damage so that maintenance can be scheduled before a potential failure. In a larger sense, the stories in this section illustrate the evolving nature of the challenges facing those who use science to help protect and serve the nation.

Researchers Contribute a Decade of Science to the SciDAC Program

CONTACT
LORI DIACHIN
DIACHIN2@LLNL.GOV

During the past decade, researchers in the Computation Directorate have made significant contributions to the DOE Office of Science Scientific Discovery through Advanced Computing (SciDAC) program. SciDAC was launched in 2001 to develop scientific computing software that would allow sustained performance of application codes on leadership-class computing facilities and to ultimately advance DOE research in areas ranging from fusion, climate modeling, and nuclear physics to turbulence and material science modeling. To accomplish this goal, the SciDAC program invested in a coordinated set of activities in three primary program elements:

Centers for Enabling Technologies: multi-institutional teams focused on developing robust algorithms that scale to hundreds of thousands of processors; methodologies for achieving high performance, portability, and interoperability of complex scientific computing software; and tools for data analysis and management.

Institutes: university-led centers of excellence intended to increase the presence of the SciDAC program in the academic community and to complement the efforts of the Centers for Enabling Technologies.

Science Applications and Scientific Application Partnerships (SAPs): scientific domain teams focused on developing scientific insight through advanced simulation. Strong, deep partnerships with computer scientists and mathematicians are fostered through SAP activities and are key to the progress of simulation-based science.

Progress in 2010

Scientists in the Computation Directorate have been integrally involved in the SciDAC program since its inception and have contributed significantly to its overall success. For example, in the last five years alone, Computation scientists collaborated with laboratories and universities on 17 SciDAC projects from the three program elements. A few of their many contributions are highlighted in this article.

Computation's Dean Williams leads the Earth System Grid (ESG) project, which enables faster, easier sharing of climate change data. ESG standardizes how data is collected and provides a common framework for the world's climate modeling centers, changing how data is accessed, shared, and compared. This technology played a critical role in the development of the Intergovernmental Panel on Climate Change (IPCC) Fourth Assessment Report. IPCC shared a Nobel Peace Prize with Al Gore for this work in 2007.

Computation's Lori Diachin leads the Interoperable Tools for Advanced Petascale Simulations project, which develops interoperable technologies such as mesh quality improvement, adaptive mesh refinement, and front tracking for unstructured mesh simulations on massively parallel computers. These technologies have been used to decrease the overall time to solution and simultaneously increase the accuracy of fusion, accelerator, and nuclear energy applications.

Award-winning LLNL software has been critical to SciDAC success. In particular, Computation researchers have earned R&D 100 awards in a number of areas, including linear solvers (*hypre*), compiler infrastructure (ROSE), language interoperability (Babel), scientific visualization (VisIt), and scientific data management (Sapphire). These tools are integral to the work product deliverables of their respective SciDAC Centers and Institutes.

As part of the TOPS (Towards Optimal Petascale Simulations) Center for Enabling Technologies, Computation researchers have furthered the development of massively parallel linear solver algorithms, which are at the core of many scientific simulation codes. LLNL's parallel linear solver library *hypre* has provided workhorse solvers to a variety of applications at LLNL, DOE (SciDAC included), and around the world. With funding from the SciDAC program, the *hypre* team made several algorithmic developments that enabled near-perfect weak scaling of their algebraic multigrid solver, BoomerAMG, up to 130,000 cores of BlueGene/L at LLNL and up to 196,000 cores of Jaguar at Oak Ridge National Laboratory. The improved solver is more than six times faster on some problems and is aiding scientific simulations in geophysical flows such as magma migration, mantle convection, and ice sheet dynamics.

Understanding the communication patterns of complex large-scale science applications is a significant challenge, largely because of the massive data collected during large-scale simulation runs. In contrast to other known

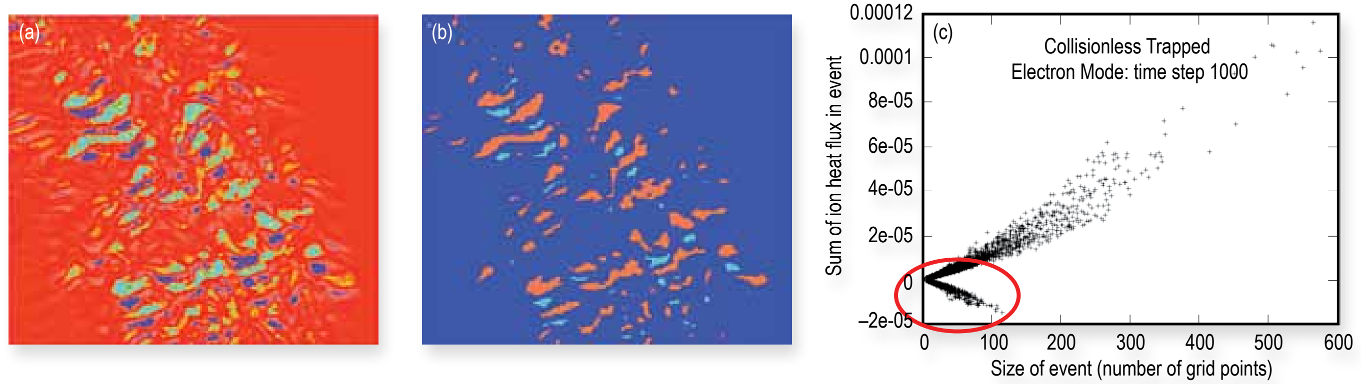

Computation's data analysis efforts identified the unexpected presence of structures with negative ion heat flux variables. Figure (a) shows a subset of the data on a two-dimensional poloidal plane from a fusion simulation showing the ion heat flux variable. Figure (b) shows the structures in (a) with positive structures in red and negative structures in blue. Figure (c) is a plot of the sum of ion heat flux in a structure vs. the size of the event and clearly shows the large number of structures with negative flux.

methods, LLNL SciDAC research conducted as part of the Performance Engineering Research Institute developed a mechanism for tracing MPI events that results in near-constant trace file sizes as problem size and/or node count increase. The compression mechanism automatically detects communication patterns, providing one of many promising areas for continued research. This work generated interest in trace-based performance analysis and demonstrated that these techniques can remain relevant, even as the computing community plans for exascale systems.

Computation researchers are analyzing fluid and particle data from fusion simulations to identify coherent structures, obtain statistics on the evolution of the structures over time, and understand the nonlinear interactions between structures. The analysis is challenging and it is difficult to identify robust algorithms because there is no definition of coherent structures, the structures vary extensively over time, and the data are spatially unstructured. The researchers' analysis of the fluid data indicated the unexpected presence of structures with negative ion heat flux that remain coherent and persist over time, suggesting that they are not due to noise. Physicists are investigating these negative structures to gain insight into energetic particle turbulence and transport, which will impact the performance of burning plasmas in ITER.

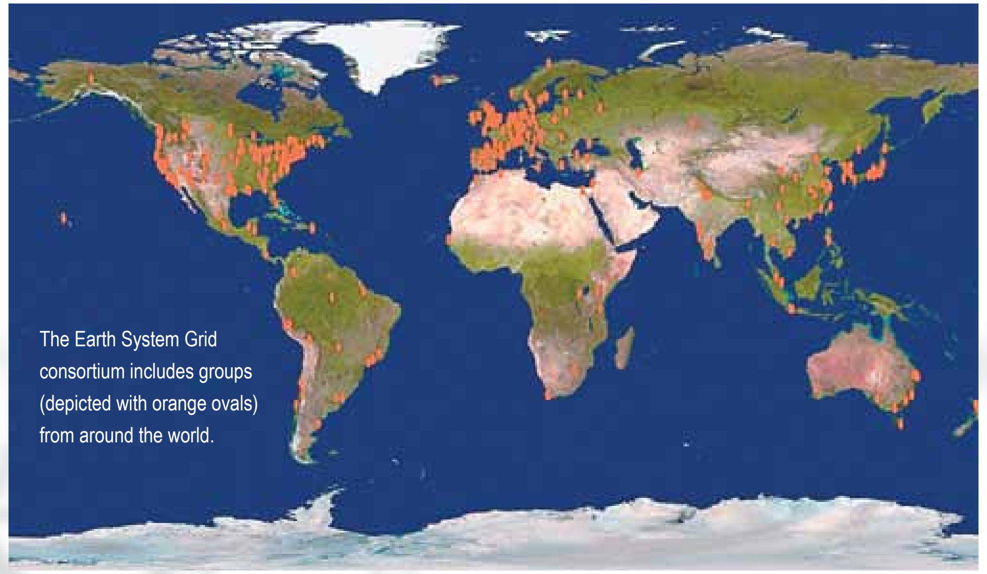

The Earth System Grid consortium includes groups (depicted with orange ovals) from around the world.

In support of the SAPs program element, Computation researchers more than doubled the maximum throughput rate of the Community Atmosphere Model on production platforms by exploiting additional parallelism, improving communication protocols, and eliminating scalability bottlenecks. The improvements allow much higher resolution and increased process representation at a cost comparable to previous implementations. The improved data is being used in the IPCC Fifth Assessment Report. In a different project, Computation researchers addressed scalability bottlenecks in solving large-scale density functional theory problems using O(N) localized orbital methods. The combination of the linear complexity algorithm with improved parallel scaling enabled calculations of problem sizes previously considered beyond reach. The method is being used on first-principles molecular dynamics simulations of biochemical reactions facilitated by an enzyme and large-scale calculations of nanostructures embedded in an electrostatic field.

The next phase of the SciDAC program will address a new decade of challenges, including the increased use of multi-core architectures and million-way parallelism. The Computation Directorate expects to continue its strong contributions in leadership, software, and research as SciDAC enters this new phase.

COMPUTATION DIRECTORATE ANNUAL REPORT 2010

Research Yields Faster, More Accurate Large-Scale Text Analysis

CONTACT
DAVID BUTTLER
DAVIDBUTTLER@LLNL.GOV

Large-scale text analysis has been a challenge for researchers, businesses, and intelligence agencies since the 1950s. Today, as the amount of information reaches new heights, organizations are seeking advanced tools and techniques to discover and present knowledge—facts, keywords, and relationships—that is otherwise locked in textual form. The overarching goal is to turn text into data for analysis. An increasing number of programs at LLNL are making large document collections available to subject matter experts and analysts, and these experts need innovative techniques for extracting relevant data from the text and structuring the information content. For example, Trinidad, a software platform supporting LLNL's counterproliferation mission, manages more than 50 million scientific abstracts and news articles spanning more than 20 years. It also contains more than 450 million unstructured records from other sources and provides instant search and retrieval of these records. But extracting the right data from these hundreds of millions of records can be like searching for the proverbial needle in a haystack. Computation is involved in several projects that are helping analysts find the most appropriate data faster and more efficiently.

Progress in 2010

Many electronic commerce companies allow customers to refine their product search by any number of attributes, such as brand, price range, or color. These "facets" are populated by fields in the company's product database. Documents are not nearly as structured; they might have source-specific metadata fields associated with them (e.g., author, institution, publication date), or they might be simple text. To create a consistent and rich user interface, a team of Computation scientists generates additional facets, such as entity annotation, domain-specific keyword lists, and automatically derived topics. This capability allows users to see how often a particular word appears in their search results. It also allows users to cross-correlate the most frequently mentioned words with the most prominent topics across an entire corpus or within their query results. The Computation team is working to enrich their set of facets by extending the concepts of natural language processing in their research.

Analysts and subject matter experts are often required to search through streams of documents; it is not unusual for analysts to need to triage and prioritize more than 1,000 messages a day. Computation researchers have stepped in to help by developing iScore, a semi-automated system that assists users in ranking documents. iScore works by combining many computed document features, including writing style, subjectivity, polarity, language models, and topic models, into a profile. The system incorporates every explicit facet associated with the documents, as well as a model of the kind of documents an analyst is seeking. Analysts interact with the profile by giving a simple "thumbs up" or "thumbs down." The system incorporates this feedback and re-ranks articles the analysts have not yet seen, thereby maximizing the probability of finding useful information in the shortest time.

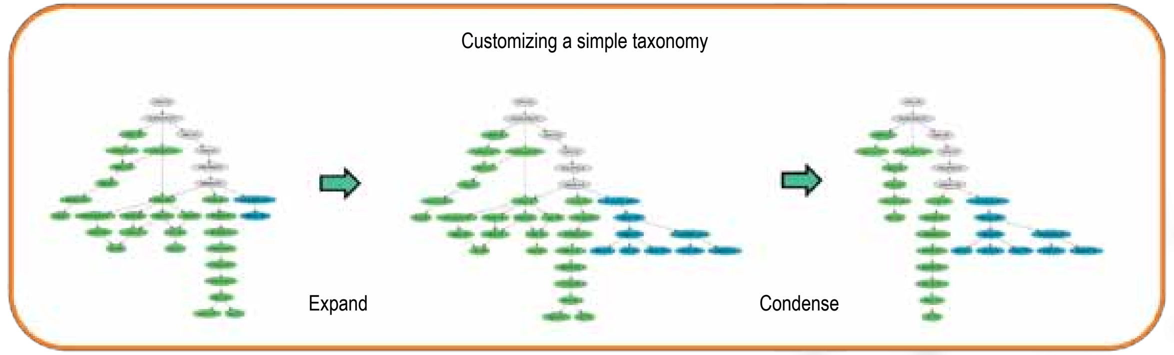

An existing concept hierarchy can be extended based on concepts found in a corpus. The concept hierarchy can then be condensed by preserving concepts that have a high word count in the document, allowing the analyst to more easily navigate the resulting ontology.

The Computation team, in conjunction with Professor Byron Gao at Texas State University, has also developed a clustering search interface that incorporates user feedback. Typical search results, as delivered by a text search engine, present results as a flat list ordered by a relevance metric. Although this method has proved to be effective when users are looking for a specific result, it is inadequate for many cases. For instance, users often approach a new corpus to find as many relevant documents as possible. The new interface helps find related documents by clustering documents and labeling the clusters with frequent phrases. In addition, users often need to organize the resulting data. The new interface allows users to edit the cluster results so that similar searches maintain the user-defined cluster edits. The interface also supports aggregating the cluster edits across users so that other users can benefit from previous searches. This capability allows an expert in one area to implicitly share their knowledge with others in their group as they improve their results.

The Computation team is also working on concept hierarchy generation. Hierarchical facets typically link nodes in a known ontology to records in a database. The main difficulty in applying this technique to documents is that most corpora do not fit in common ontologies; developing custom ontologies is too expensive given the rapidly changing set of relevant documents and the short-term nature of tasks. The team is developing techniques to automatically locate concepts in a corpus and then build relationships between those concepts and a more general concept hierarchy (e.g., WordNet). The resulting ontology is collapsed based on the word usage in the target corpus. The result is a corpus-specific extension to familiar hierarchies and more directed browsing.

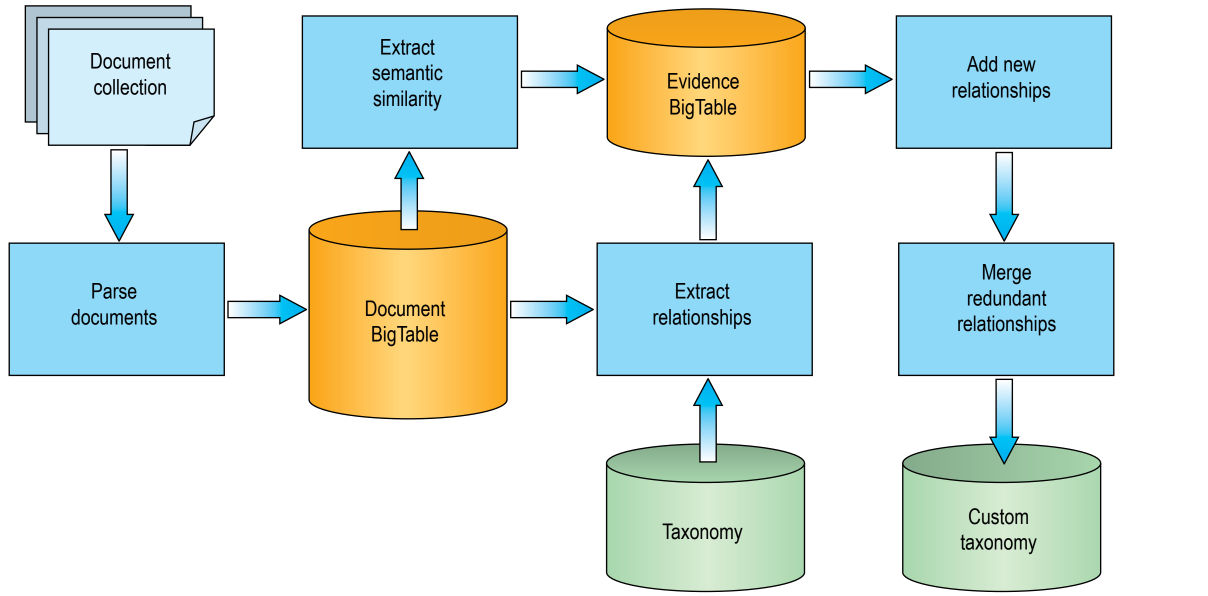

This sample processing pipeline for customizing taxonomies shows the steps to process a corpus, or document collection, to produce a custom concept hierarchy related to that document collection. The processing is based on MapReduce technology and can therefore scale with the number of processors or disks in the system.

Analysts often compile lists of terms they want to use to enrich their search results. These lists are often subject-specific jargon or acronyms they expect to find in relevant articles. Unfortunately, many acronyms and terms in the English language are overloaded with multiple meanings—most of which are irrelevant to the analyst's subject area. Word-sense disambiguation examines the context in which a word is used and statistically matches the word to a known or novel sense. The Computation team is researching ways to improve the accuracy of automated disambiguation systems to reduce incorrect matches while retaining a very low miss rate.

The team is also developing techniques that address topic modeling. Latent Dirichlet allocation-based topic modeling has yielded exciting research in statistical natural language processing for several years. Unfortunately, this research has resulted in very few functional user interfaces, mostly due to the cost of presenting probability distributions over words in coherent ways that allow users to improve their search experience. The Computation team's techniques automatically select high-quality topics, measured by their point-wise mutual information over explanatory texts. The team plans to demonstrate how to incorporate these results into a search interface to improve either precision or recall, depending on a user's need. The team is also developing techniques to reduce the cognitive load of presenting topics to users by generating more concise and descriptive topic labels.

RETURN TO TOC

COMPUTATION DIRECTORATE ANNUAL REPORT 2010

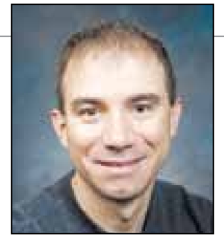

High-Order Finite Elements Improve Lagrangian Hydrodynamics Simulations

CONTACTS
TZANIO KOLEV
TZANIO@LLNL.GOV
VESELIN DOBREV
DOBREV@LLNL.GOV

Numerical simulations of shock hydrodynamics phenomena are critical to many application areas at LLNL, including the National Ignition Facility (NIF) and stockpile stewardship. Many of the large-scale software capabilities developed at LLNL for these problems employ the Arbitrary Lagrangian–Eulerian (ALE) framework. The Lagrangian step, which is at the heart of all ALE algorithms, has a number of long-standing discretization issues that can limit the predictive capability of computer simulations. These issues include symmetry breaking and low-order accuracy on general grids, lack of total energy conservation in axisymmetric problems, and mesh imprinting due to hourglass-mode instabilities and artificial viscosity treatment.

To address the above deficiencies, Computation research scientists collaborated with Robert Rieben from the Weapons and Complex Integration organization to develop a new Lagrangian discretization framework based on high-order finite-element field representations, curvilinear zone geometries, tensor artificial viscosity, and high-order time stepping algorithms. Although this approach can be viewed as a high-order generalization of traditional schemes, it has a number of advantages, such as increased accuracy and robustness on unstructured grids and significant improvement in symmetry preservation for radial flows. The new discretization methods were implemented in the recently developed research code BLAST. The results from an extensive set of two-dimensional (2D), three-dimensional (3D), and axisymmetric tests demonstrate that high-order curvilinear finite elements enable higher quality Lagrangian simulations than previously possible.

Progress in 2010

The mathematical description of shock hydrodynamics phenomena is given by the Euler equations, which express the conservation of momentum, mass, and total energy. In the Lagrangian step of the commonly used ALE approach, these equations are written in a material frame and are traditionally discretized with staggered grid hydrodynamics (SGH) schemes where kinematic unknowns, such as the velocity and the position of the fluid particles, are located at the mesh vertices, while the thermodynamic unknowns, which include the density, pressure, and internal energy, are cell centered. In contrast, the newly developed finite-element framework uses additional degrees of freedom to describe these variables as functions that vary inside each zone. For example, the "discontinuous Taylor–Hood" approach in 2D uses nine velocities and four energies per zone. A high-order kinematic description is beneficial as it enables curvilinear zone geometries, allowing for better approximation of the mesh curvature that develops with the material motion. In addition, high-order functions allow for better approximation of the unknown fields in smooth regions and capture subzone dynamics. From the general high-order framework, one can derive many of the classical SGH schemes by choosing low-order kinematic and thermodynamic bases.

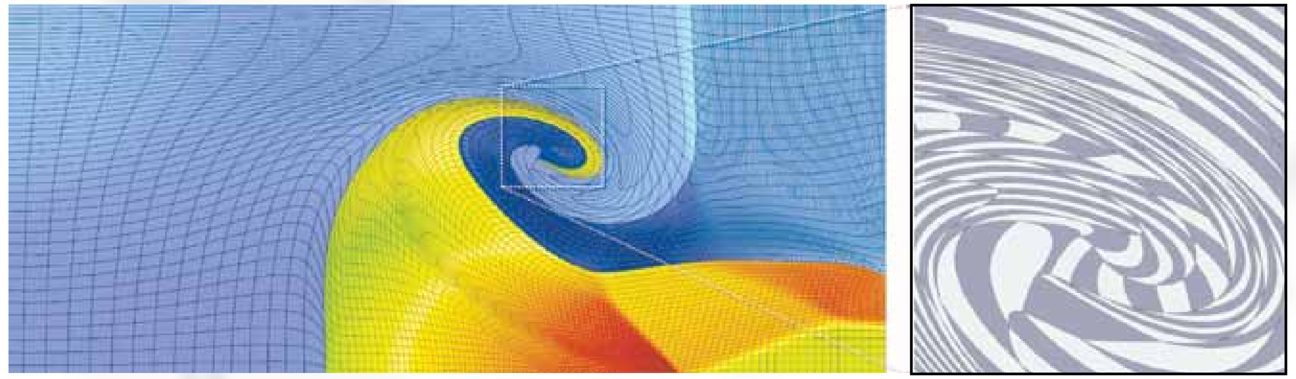

The new curvilinear finite-element discretization technique expands the applicability of pure Lagrangian methods to complicated flows producing highly deformed and stretched zones. Shown is the log-scale density of a two-dimensional multi-material Riemann problem describing a shock triple-point interaction and a close-up of the curvilinear mesh around the point of maximum vorticity. The resulting deformed mesh cannot be accurately represented using zones with straight edges.

The functional point of view leads to new insights for the discretization of the Euler equations. One important example is the concept of strong mass conservation, which defines the density as a function based on the initial density and the current deformation. This concept was essential for the robustness and efficiency of the new finite-element approach. One of its applications is in the variational formulations of the momentum and energy conservation equations with a high-order discontinuous basis for the internal energy field. The resulting linear equations involve a global "velocity mass matrix" and a set of small local "energy mass matrices" on each zone. By strong mass conservation, these matrices are independent of time and are, therefore, only computed once at the beginning of the computation.

The forces appearing on the right side of the momentum equation are computed through a high-order generalization of the SGH "corner force" concept. These generalized corner forces can be represented as local rectangular matrices on each zone, and their evaluation forms the main computational kernel of the method. This locally floating-point-operation-intensive computational kernel is well suited for modern heterogeneous computing architectures where the emerging trend is that floating-point operations are much less expensive than data motion in and out of local memory cache.

While the finite-element framework is general enough to handle 2D and 3D problems, in practice 3D problems with axial symmetry can be solved much more efficiently through a reduction to a 2D meridian cut in cylindrical coordinates. Traditionally, Lagrangian methods for such axisymmetric problems have used the "area-weighted" approach, where the momentum equation is solved in 2D planar coordinates using the "area masses" at nodes while the internal energy equation is solved over the real volumes. This approach does not conserve total energy and can often lead to spurious symmetry breaking in the internal energy field near the axis of rotation. In contrast, the proposed extension of the high-order finite-element framework to these problems conserves total energy by construction. The new method was derived by a careful reduction of the 3D axisymmetric problem to a 2D variational form in the meridian cut. Unlike the SGH area-weighted schemes, the finite-element approach leads to a rescaled momentum conservation equation, which also includes new terms in the pressure gradient and artificial viscosity forces.

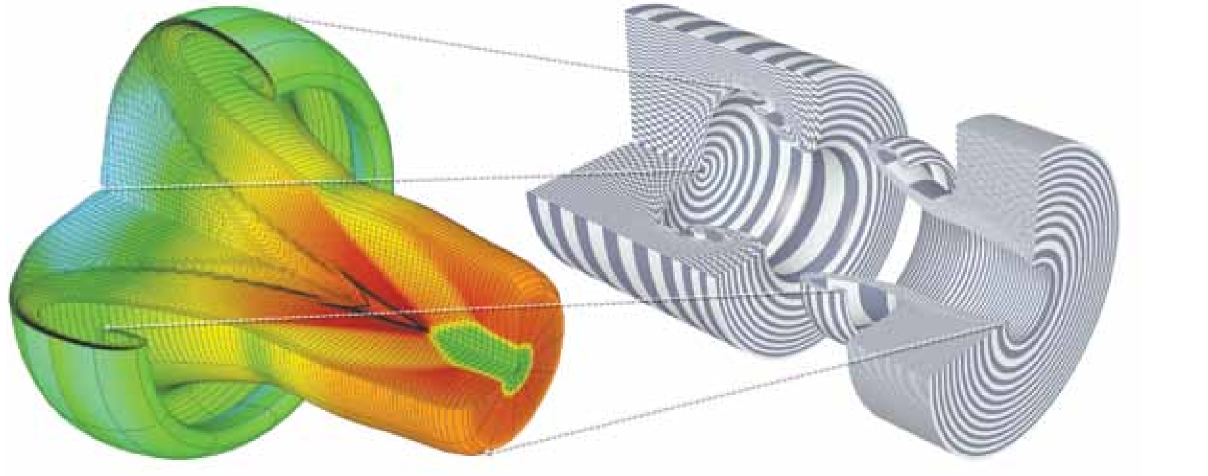

The extension of the curvilinear finite-element framework to axisymmetric problems leads to one of the few known approaches that support both symmetry preservation and exact total energy conservation in cylindrical geometry. Shown are the revolved three-dimensional density and the curvilinear mesh computed from the axisymmetric version of the shock triple-point interaction problem in the figure on page 12. The axisymmetric method maintains robust performance similar to the two -dimensional case and does not produce artifacts near the axis of rotation.

As in Cartesian coordinates, the use of high-order basis function expansions, obtained via a high-order mapping from a standard reference element, enables the handling of curvilinear zone geometry and higher order approximations for the fields within a zone. The axisymmetric version can be viewed as a relatively simple modification of the 2D case, where the velocity and energy mass matrices have a radial weighting function, while the generalized corner forces also include the terms mentioned above. Based on the concept of strong mass conservation in cylindrical coordinates, these matrices remain constant in time and lead to exact total energy conservation on a semi-discrete level, as in the Cartesian case.

The implementation of the high-order curvilinear finite-element discretization methods in BLAST was verified using an extensive set of problems designed to exercise various important aspects of Lagrangian hydrodynamics. The results demonstrate the advantages that the high-order approach can provide, including increased accuracy and robustness on unstructured grids, significant improvement in symmetry preservation for radial flows, exact (up to machine precision) energy conservation, high-order convergence for smooth problems, elimination of the need for ad-hoc hourglass filters, sharper resolution of shock including the ability to represent shocks within a single zone, and a substantial reduction in mesh imprinting for shock waves not aligned with the computational mesh.

RETURN TO TOC

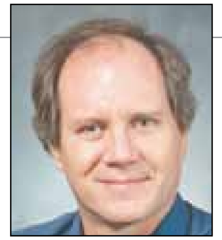

Compiler Promotes Novel Uses of Automated Code Analysis and Transformation

CONTACTS
DANIEL QUINLAN
DQUINLAN@LLNL.GOV
CHUNHUA LIAO
LIAO6@LLNL.GOV

ROSE is an open source compiler project at LLNL that is contributed to and used by research groups worldwide. Compilers take source code written by humans (typically) and translate it to code that can be executed by machines. As a compiler infrastructure, ROSE supports the languages used within DOE applications, specifically C, C++, Fortran (77, 90/95, 2003, and 2008), Open Multi-Processing (OpenMP), Unified Parallel C (UPC), Compute Unified Device Architecture (CUDA), and Open Computing Language (OpenCL). ROSE is unique in that it is a source-to-source compiler addressing the custom requirements of application groups to build automated tools that operate directly on source code. A vendor compiler is then used on the ROSE-generated source code to create the machine code. This approach makes ROSE portable and more broadly targeted to custom or application-specific optimizations. One of the goals of the project is to support the custom analysis and optimization requirements of scientific applications within DOE. Custom transformations (e.g., optimizations) can be automated on large-scale source code to focus on customized uses of hardware features or optimizations that are high level or specific to a particular application or parallel programming model. Since its inception more than 12 years ago, ROSE has continuously gained popularity, and research groups internal and external to DOE are still finding new ways to use it to meet their compiler needs.

Many laboratories and universities are using ROSE to help prepare for and analyze evolving computing requirements. For example, exascale architecture design teams at Argonne National Laboratory and LLNL used ROSE to characterize the behavior of their laboratory codes on their planned exascale systems. This approach to building new computing architectures, where both the hardware and the software are studied to understand how to build improved software and hardware recursively, is called codesign. Lawrence Berkeley National Laboratory, Sandia National Laboratories, University of Illinois at Urbana-Champaign, and UC San Diego used ROSE to support parameterized processor simulators within hardware/software codesign to define custom analysis and transformations. Their goal was to exploit novel exascale architecture designs (e.g., specialized instructions for scientific applications, active control of power usage by applications, compiler optimizations). Active power management and design will be critical in DOE's exascale computer architectures. These architectures will need to be more power efficient (fewer watts/floating-point operations) to be feasible to build and maintain. LLNL's work with ROSE permits early experience with future design features to explore both the hardware features and how to write the software and/or design the compiler optimizations to use the hardware features; this is the essence of hardware/software codesign. Future systems architectures might also be heterogeneous.

Los Alamos National Laboratory scientist Craig Rasmussen used ROSE to demonstrate a 40× performance improvement through the automated translation of Fortran 90 array constructs (source code) to OpenCL for application kernels deemed appropriate for use on GPUs. ROSE is also being used to support the analysis of traces generated from the execution of programs and cyber security applications.

The optimization of applications is critical within DOE. The ROSE project team, in collaboration with researchers from the DOE SciDAC Performance Engineering Research Institute, has dedicated significant effort to optimizing performance using a technique called auto-tuning. The team has specifically focused on whole-program auto-tuning to support large-scale DOE applications. In auto-tuning, many different versions of parameterized optimizations are implemented and tested off-line, and the best version of the optimization is selected for the final compilation of an application. Auto-tuning permits an organized approach to the optimization of applications where internal compiler cost-modeling is intractable. This technique is especially useful when the performance is sensitive to minor perturbations in floating-point pipeline length, cache-line size, memory access stride, or any other case where the architecture is complex and not easily modeled within the compiler. Unfortunately, computer processor architectures are growing in complexity, and it is possible that auto-tuning will not become simpler in exascale architectures. Exascale architectures could be routinely optimized using this technology if research can make it more efficient for large-scale applications.

Progress in 2010

In 2010, the ROSE project team focused on supporting several large-scale DOE Fortran 90 applications and defining customized analysis and transformations, adding CUDA and OpenCL language support to ROSE (for graphics processing unit, or GPU, programming), researching parallel programming models (e.g., MPI and OpenMP), and completing the auto-parallelization of serial programs to support multi-core processors. The team developed a full OpenMP 3.0 compiler using ROSE, which provides researchers from around the world a user-friendly open source OpenMP platform with which to explore extensions supporting exascale computing.

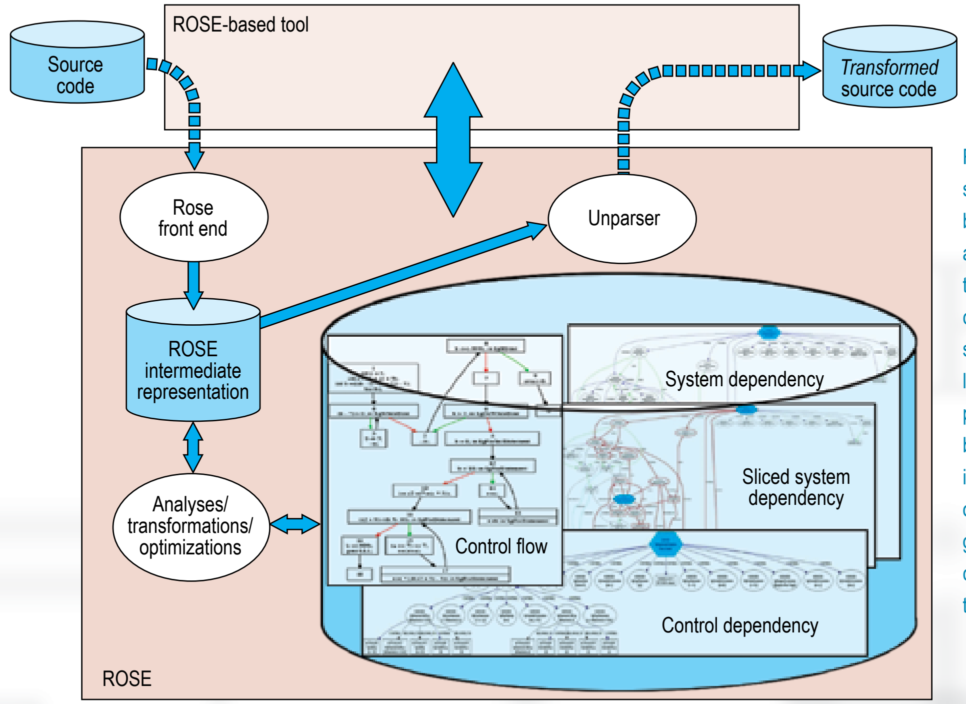

ROSE defines an open source infrastructure for building custom tools for arbitrary analysis and transformation of source code. Analyses include searching for parallelism, looking for performance problems, and finding bugs. Transformations include instrumentation, optimizations, and generation of entirely new codes with parts specific to GPUs.

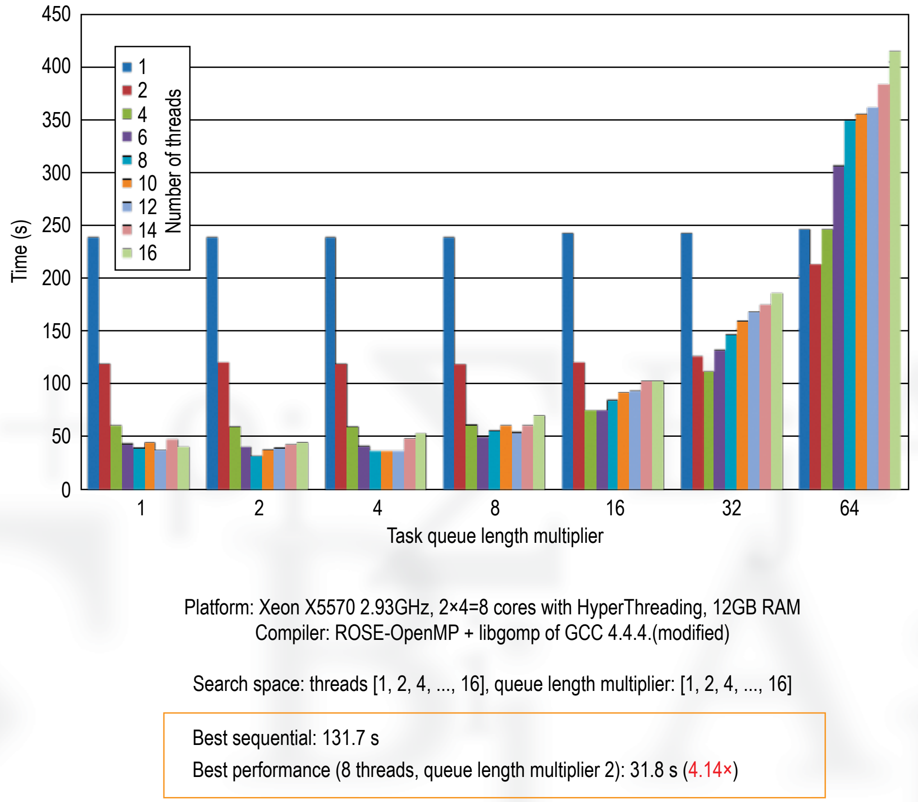

The performance of an application with task parallelism is measured while varying the number of threads and the task queue length. The optimal length of the queue for this section of code is explored explicitly using auto-tuning, and it is then used in the final execution (using the ROSE OpenMP transformation). The result is a 4.14× improvement in performance compared to the parallel execution using OpenMP with the default runtime parameter settings.

COMPUTATION DIRECTORATE ANNUAL REPORT 2010


Persistics Delivers Scalable Computation for Wide-Area Persistent Video Surveillance

CONTACTS
JONATHAN COHEN
JCOHEN@LLNL.GOV
MARK DUCHAINEAU
DUCHAINEAU1@LLNL.GOV

Wide-area persistent surveillance sensors gather aerial imagery covering hundreds of square kilometers on the ground for durations of hours, days, or more. These assets provide critical capability in such areas as tracking terrorist threats and exposing proliferation networks. However, the data rates produced by the largest of these sensors may exceed 1 terabyte per minute. The sheer volume of data is an obstacle to transmission, analysis, and ultimately the production of actionable intelligence.

The Persistics project is developing scalable algorithms running on high performance, heterogeneous computing architectures to address these challenges. The algorithms address issues such as sensor nonuniformity correction, vignette correction, seamless "stitching" of multi-sensor images, static and dynamic bias-gain adjustment, stabilization, geo-registration, parallax removal, foreground/background modeling, tracking, and, most importantly, compression. Looking forward, Persistics is moving further up the semantic chain, exploring issues such as automatic road network extraction, modeling of normal behavior as a function of location, and anomaly detection. The project was initially funded internally through the VidCharts LDRD but has since transitioned to the Department of Defense and received additional support from DOE through the nonproliferation program.

In November 2010, one instance of the 12-node prototype ground station was delivered to the sponsor's site, where it will undergo extensive testing on new data sets, enabling refinement of the algorithms to improve robustness during the coming months. A second instance of the 12-node cluster remains at LLNL, where it will serve as the baseline for new research and development in 2011. Ongoing efforts include scaling the processed video data rates by a factor of 30, reducing power requirements, developing three-dimensional extraction capabilities, incorporating more sophisticated tracking technologies, and providing a higher level of semantic analysis to assist human intelligence analysts.

Progress in 2010

In 2009, Persistics demonstrated a novel algorithm for high-quality video compression of aerial video with compression rates exceeding 1,000-to-1, more than an order of magnitude improvement over what was previously considered state of the art. This fundamental advance is possible because of the persistent nature of the input imagery—the same area on the ground is imaged repeatedly. The Persistics compression algorithm analyzes the imagery to produce a high-quality background model, which changes only with the time of day and the camera's viewing direction. Pixels that differ from this background model are stored in a separate foreground channel, capturing items such as vehicles, people, dust, and clouds. The Persistics compression pipeline essentially preconditions the video data using many layers of corrections, including sub-pixel stabilization and geo-registration, parallax removal, and static and dynamic intensity corrections to minimize residuals before passing the data through a standard H.264 video encoder.

This year, the Persistics project transformed its demonstration software into a full-scale, prototype, ground-based video processing system. The two-rack, 24-kilowatt cluster consists of 12 compute nodes, each containing two quad-core Intel processors and seven NVIDIA Fermi C2050 graphics processing units (GPUs), with InfiniBand interconnect. The Persistics software pipeline is distributed over the compute nodes and processors using a dataflow graph abstraction, where elements of work to be performed and data are represented as nodes in a graph, and edges connect the data and functions to indicate dependency relationships. Processing is performed using a multi-resolution hierarchy, with pixel-level parallelism on each of the 480-core GPUs, and tiles of geo-space distributed across nodes and processors.

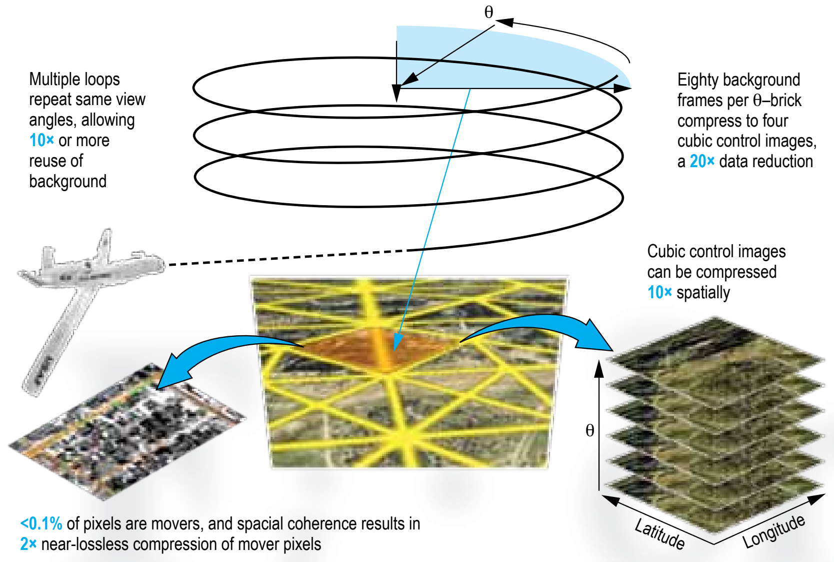

Persistent flight loop images are diced into latitude-longitude-theta bricks, decomposed into background and foreground, and compressed 1,000 times or more.

RETURN TO TOC

Data Mining Aids Optics Inspection at the National Ignition Facility

CONTACT
GHALEB ABDULLA
ABDULLA1@LLNL.GOV

The Final Optics Damage Inspection (FODI) camera system is inserted into the center of the National Ignition Facility (NIF) target chamber after a laser shot to acquire images of the final optics for all 192 beamlines. The images are then processed using custom image processing and analysis software, referred to as the Optics Inspection (OI). OI computes several measurements, such as area, diameter, and intensity, for each beamline site. This information helps scientists track sites along the beamlines that were damaged in a previous shot. The damage sites are recorded and monitored and then scheduled for maintenance before they reach a certain size. Artifacts, such as hardware reflections, complicate this task because they are difficult to distinguish from legitimate damage sites. Computation researchers are using data mining techniques to help clean artifacts from the data and characterize and understand the nature of damage growth behavior.

Progress in 2010

Computation researchers began the process of cleaning hardware reflections from the optics data (i.e., filtering out false-positive results) by retrieving images of potential damage sites from a NIF image database, along with their classifications as defined by NIF experts. Because several experts contributed labels for the data in relatively unconstrained ways and at different times, no standardized categories were used to describe the sites they examined. The Computation team, however, successfully grouped the experts' comments into three standard classes: hardware reflection, not a reflection (i.e., damage site), and unknown.

Researchers compiled the data from the first two categories (hardware reflections and damage sites) into one data set that they used to train a classifier. Features, such as area in pixels, x-location, y-location, and long axis, were obtained using the OI software package. The team then entered this data into a classification tool called Avatar, which was developed by Sandia National Laboratories and the University of Florida.

Using Avatar, the team created an ensemble of decision trees to design a model that they then tested using tenfold cross validation. The model was more than 99% accurate, largely due to the high-quality feature set the researchers had gathered when they initially separated the data into distinct classes. The team then used the model and training data to attempt to separate data from the third group (initially classified as "unknowns") into either the hardware reflection or damage site category. The classifier results were shown to the experts to verify the model prediction; the results matched the experts' opinions 99% of the time. The data set and the classifier model are now being used to filter hardware reflections in real time from the data collected by the FODI camera.

The Computation research team is continuing to help NIF experts understand the nature of the damage sites' growth behavior, including identifying the surfaces that are more prone to damage, understanding how the morphology of the site affects the growth rate, and defining the parameters that govern the growth behavior. The goal is to increase the current prediction accuracy of future damage sites.

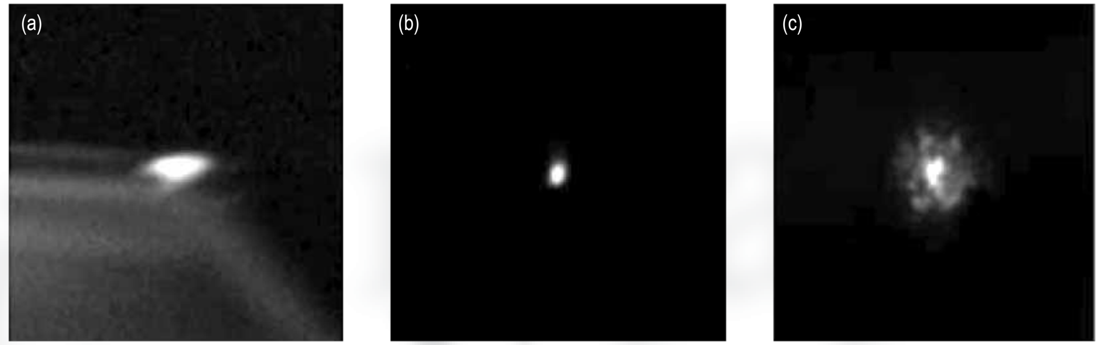

Examples of images from the National Ignition Facility's Final Optics Damage Inspection camera system were classified by the Computation research team as: a) hardware reflection, b) damage site, and c) unknown.

COMPUTATION DIRECTORATE ANNUAL REPORT 2010

Eigensolvers Help Glean Information from Extremely Large-Scale Graphs

CONTACT
VAN EMDEN HENSON
HENSON5@LLNL.GOV

Many important problems are modeled using graphs, or collections of vertices and edges connecting the vertices. In the cyber security and intelligence domains, people or computers are often represented by the vertices, while relationships connecting them, such as text messages, phone conversations, e-mail exchanges, or data transmissions, are represented by edges. Graph algorithms enable the discovery of relationships among vertices such as communication patterns, data flow, connectivity, communities, and associations. The graphs can be represented with matrices; linear algebra can be used to model and analyze the graphs. Spectral graph algorithms are one method of analysis that has wide-ranging applications, including ranking the importance of vertices, graph partitioning, or community identification; computing "distances" on the graph; and determining the ease of communication among vertices. Graphs can grow to extraordinary sizes, sometimes O(10^{10}) vertices, often evolve rapidly in real time, and typically have properties that make analysis extremely difficult.

Progress in 2010

Spectral methods require the computation of eigenvalues and eigenvectors of the graph matrices, a difficult and expensive process even for small matrices. Large-scale eigenvalue methods have been studied and perfected for most applied physics problems and run effectively on LLNL's massively parallel distributed-memory computers. However, on power-law scale-free graphs, the classic algorithms are unscalable and frequently fail entirely. The Eigensolvers LDRD project was established in 2010 to invent scalable methods for computing the eigenvalues and eigenvectors (eigenpairs) for these large-scale graphs. Research in the first year proceeded on two main fronts.

First, researchers used scalability tests and analysis to discover that a major impediment to parallel scalability of classic spectral algorithms is the standard matrix-vector multiplication (matvec). This principal step occurs repeatedly in all eigenvalue methods and is unscalable for power-law scale-free graphs because these graphs feature long-distance connections and individual vertices with many connections (hubs) that must communicate with many other vertices. Such vertices require excessive inter-processor communication at each matvec. Hence, the research team initiated a multi-pronged effort to improve matvec scalability. They made significant progress using alternative storage and communication schemes; and, while theoretical considerations limit the total improvement possible with this approach, the team is now focused on maximizing that improvement and determining the effects of these changes on the overall spectral method. The team also discovered a method of "disaggregating" hubs to produce an alternative graph whose eigenpairs approximate those of the original graph to arbitrary accuracy and have better communication properties. This discovery opens the door to a host of algorithms that could evolve into significantly improved methods.

The team's second effort was to apply Computation's deep expertise in algebraic multigrid (AMG) methods to find the graph eigenpairs. AMG is normally applied to graphs that are neither power-law or scale-free, so devising the scalable graph hierarchies that preserve these properties on all levels was both crucial and challenging. The team constructed two serial AMG-based eigensolvers, both employing forms of aggregation coarsening. The first eigensolver was specifically intended for the vertex-ranking problem, and it has proven effective and scalable on a specific category of graphs. The team is working to extend it to more general graphs. The second eigensolver, a sophisticated blend of AMG and classic Lanczos methods, was designed for the more difficult "commute-time" problem. The team is testing a very preliminary implementation. Parallelization of both algorithms will follow.

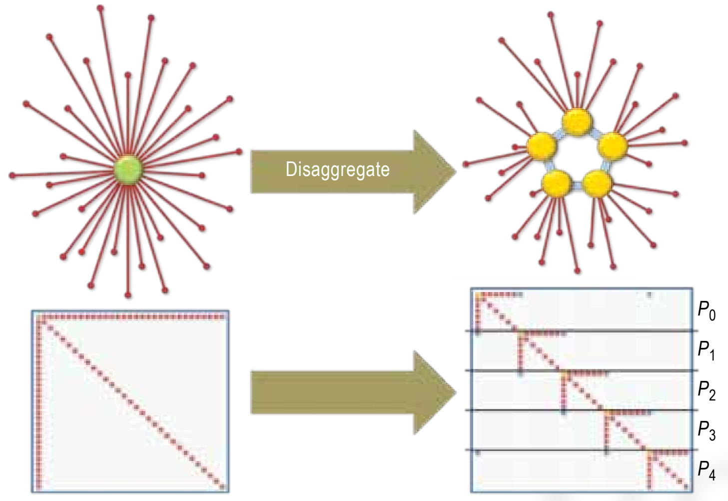

A graph can be represented as a matrix, where each node is a row and column and an entry indicates that the nodes are connected. Disaggregation (top) breaks a graph "hub" into connected sub-vertices and replaces the nonparallelizable graph matrix with a parallelizable matrix (bottom). Each portion of the matrix can be sent to different processors (P0, P1, etc.) for parallel processing. The resulting matrix has eigenvalues that can be made arbitrarily close to those of the original matrix.

High Performance Computing Propels Scientific Innovation and U.S. Competitiveness

TERRI QUINN
QUINN1@LLNL.GOV

The Computation Directorate at Lawrence Livermore National Laboratory (LLNL) is committed to promoting wider usage of and providing access to high performance computing (HPC) resources. In 2010, the directorate sought opportunities to expand the user base, both internal and external to the Laboratory, and to grow the directorate's already substantial HPC capabilities to keep pace with increasing demand.

This year, Computation established new partnerships and deployed new services. Directorate staff designed, installed, and now operate an HPC cluster for exclusive use by LLNL's Global Security organization. Global Security saved time and money because of Computation's ability to quickly install and operate the new cluster. Computation also created its first grid service with the deployment of the Green Linux Compute Cluster (GLCC). GLCC is being used by scientists around the globe to process data from the ALICE (A Large Ion Collider Experiment) detector, one of four major detectors on the Large Hadron Collider particle accelerator at CERN (European Organization for Nuclear Research). Computation is pursuing additional partnership opportunities by establishing an unclassified, unrestricted, "open" HPC center in the Livermore Valley Open Campus. An open HPC presence will allow the Laboratory to partner with industry to develop and disseminate HPC capabilities. The potential for HPC to deliver significant economic benefit to the U.S. in the next two decades is enormous.

Computation staff completed an unprecedented number of system installations and upgrades in 2010. By deploying five new systems in the last six months of 2010, the directorate increased its core computing capability by 25%. These systems bring the Laboratory's HPC peak capability to 2.5 quadrillion floating-point operations per second (petaFLOP/s). The delivery of the Sequoia system in early 2012 and the eventuality of exascale computing present Computation with several challenges. Sequoia, a 20-petaFLOP/s system, will have 1.6 petabytes of memory and 1.6M compute cores. This formidable system represents a massive increase in compute power for LLNL. In preparation, Computation upgraded the electrical infrastructure for HPC to carry up to 30 megawatts of electrical power. In addition, the directorate doubled the number of storage tape slots in anticipation of a substantial increase in simulation data. The new slots, along with a fivefold increase in tape capacity, are an order-of-magnitude increase in storage capacity. Computation experts are also supporting projects, such as the Sequoia Application Preparation project, to assist simulation developers as they prepare for Sequoia's 1.6M cores. With Computation's help, codes teams are making effective use of early delivery Sequoia hardware and simulators.

Computation is dedicated to the idea that HPC propels scientific innovation and U.S. competitiveness. To this end, the directorate's efforts are and will continue to be focused on paying careful attention to delivering superior HPC today while growing and nurturing an expanded constituency of HPC users.

Researchers Evaluate Programming Models for Exascale Systems

CONTACT
EVI DUBE
DUBE1@LLNL.GOV

As supercomputing moves toward exascale systems that have more than a thousand times the computing power of current supercomputers, node architectures will change significantly. Central processing unit (CPU) core counts on nodes will increase by at least an order of magnitude. In addition, heterogeneous architectures will become more commonplace and will include graphics processing units (GPU) or field-programmable gate arrays to enhance computing power. These changes will affect Laboratory code teams as they scale and prepare to run their applications.

Novel programming models, such as Unified Parallel C (UPC), Compute Unified Device Architecture (CUDA), and the Open Computing Language (OpenCL), may optimize on-node parallelism in these new architectures. Computation researchers examined these models to determine their suitability for LLNL scientific application codes. The team ported codes to the new models, measured performance, and characterized the results.

Progress in 2010

For exascale computing, the main programming challenges and biggest design changes are expected to occur on-node rather than across nodes. The total number of nodes will not increase dramatically, and as a result, Message Passing Interface—a communications protocol that allows processes to communicate with one another by sending and receiving messages—provides one option for inter-node communication on an exascale system. Several programming models have been proposed for computation tasks that occur on a single node: the Open Multi-Processing (OpenMP) standard, UPC, Co-Array Fortran, and GPU-centric models such as CUDA or OpenCL, can all be used to achieve intra-node parallelism. UPC or Co-Array Fortran can be extended to the global level.

Researchers conducted a study of programming models to resolve how application teams could effectively use the Laboratory's future exascale environment. The study's core objectives were to determine: a) relative ease in acquiring the necessary programming skills; b) relative ease in porting code; c) robustness of the development environment and tools associated with the language model; and d) performance of applications coded with the new model.

UPC, CUDA, and OpenCL were selected as the target language models for the study. UPC can either be used as an alternative to OpenMP for achieving on-node parallelism, or for attaining intra- and inter-node parallelism. CUDA and OpenCL allow code to be parallelized using GPUs. This characteristic is of particular interest because GPUs are one possible means for achieving the high levels of intra-node parallelism needed for exascale systems.

The Computation team ported an application, known as CLOMP, to each of the target language models and then compared the results to those of the same application parallelized with OpenMP. Performance characteristics were measured for a separate Laboratory application (LIP/LEOS) ported only to UPC. The team measured the time it took to learn the language, as well as the application's performance, the programming environment, and in some cases, the computing platform. The results are depicted in the table below.

Total	Application	Language
37 days	LIP/LEOS	UPC
31 days	CLOMP	UPC
25–30 days* * note: this is still not running	CLOMP	CUDA
35 days	CLOMP	OpenCL

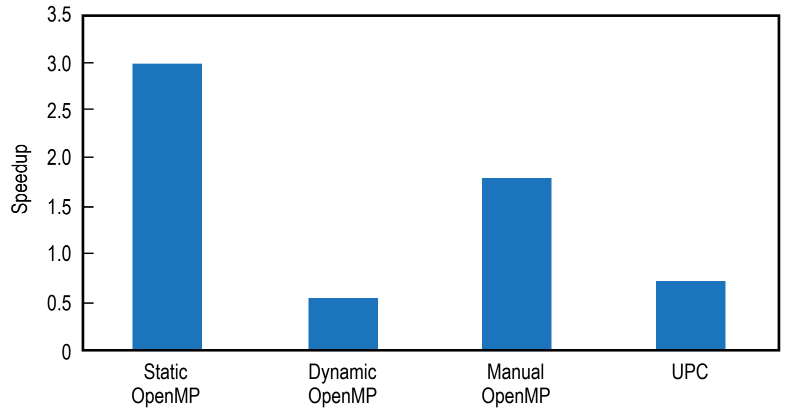

Speedup results achieved on the 16-way Hera node relative to OpenMP serial reference case for a given target input show that UPC performance was significantly worse than the serial reference, as seen by a UPC speedup value of less than 1.

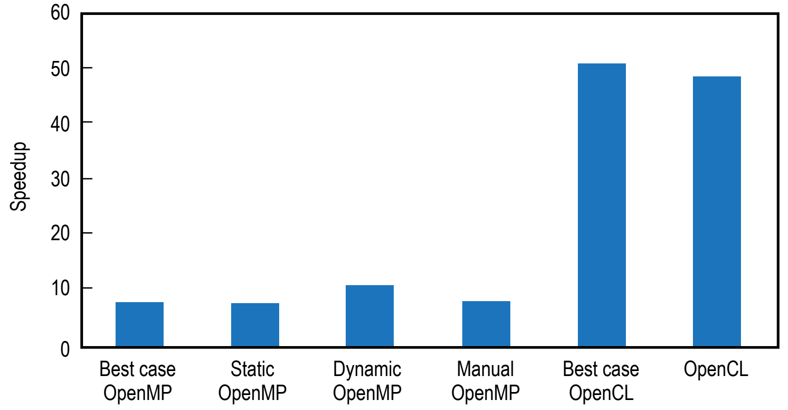

OpenCL (GPU) speedup on the 12-way Edgelet node is shown relative to the OpenMP (CPU) serial reference case for floating-point-operations (FLOP)-bound input (increased FLOP per iteration). The CPU cases are all comparable, but the GPU cases show significant speedup.

Learning the nuances of each language was challenging. In addition, debugging tools were limited, resulting in the need to use printf functions. For example, UPC-shared variables are always global, and must therefore be declared at the beginning of the program. Because UPC launches all threads at application start up, the developer is forced to examine every code file to determine what must be serialized under UPC, rather than what can be parallelized with work sharing. As another example, the lack of pointer support in OpenCL required a restructuring of the main data types (linked lists) into linear arrays suitable for transferring memory copies to and from an OpenCL device.

In the study, speedup of each ported code was compared to the OpenMP serial reference case for the given input problem. The performance runs were conducted on the Hera cluster, which has four quad-core Opteron 8356 2.3-gigaherz CPUs totaling 16 cores per node and 32 gigabytes of memory. Over a variety of problem sizes, the raw time to run parallel regions of the UPC port of CLOMP was faster than dynamic OpenMP scheduled loops. When shared memory was allocated by each worker thread, the UPC port was slower than manual or statically scheduled OpenMP loops. Combined with the generally slower serial regions of the code under UPC, the present UPC port is significantly slower than an OpenMP port with an intelligent allocation strategy.

Overall speedup of the UPC port was compared to that of the OpenMP serial reference case for a given target input that represented a small memory footprint problem. This case involved 64 partitions and 100 zones per partition, which translates to 64 independent linked lists with 100 elements in each list. Results showed that UPC performance was significantly worse than the serial reference case.

The OpenCL port of CLOMP results were not impressive, except when the problem was constrained by floating-point operations. In most cases, the OpenCL version was in fact slower than the serial OpenCL version with no parallelization. This result was expected because of the many memory buffer copies to and from the compute device. An OpenCL version that minimized these copies was not much better than the serial adaptation. Speedup may have been improved if the kernel code had been optimized.

When the amount of floating-point operations per iteration was increased by a factor of 10, an overall decrease in iterations was achieved (since the timescale was not altered). Results of this test showed significant speedup because the copy operations for OpenCL and GPU memory buffers were mitigated, illustrating the sheer amount of computations per iteration plays to the strengths of graphics processors and the raw floating-point-operation capability of the GPU.

COMPUTATION DIRECTORATE ANNUAL REPORT 2010

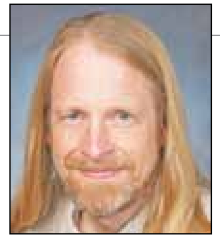

Zetabyte File System Improves Lustre Efficiency

CONTACTS
BRIAN BEHLENDORF
BEHLENDORF1@LLNL.GOV
JIM GARLICK
GARLICK1@LLNL.GOV

Livermore Computing (LC) builds high-performance Lustre file systems to store application data sets and checkpoint/restart files. As compute platforms evolve to petascale and data sizes grow proportionally, Lustre bandwidths must be increased to keep computation resources working efficiently. The Sequoia IBM BlueGene/Q platform, due in 2011, will require an estimated minimum input/output (I/O) bandwidth of 500 gigabytes per second (GB/s) from its Lustre file system, with a stretch goal of 1 terabyte per second. If this system were built using today's production components, it would cost $43M to reach the minimum bandwidth.

An effort is underway to replace Lustre's back-end object store with the Zetabyte File System (ZFS), which will improve Lustre's data-moving efficiency and greatly reduce the cost of the Sequoia file system. This project, conducted in collaboration with Oracle's Lustre group, reached a milestone in September 2010 when LC demonstrated an unoptimized prototype that, if built out to Sequoia scale, would meet the minimum bandwidth requirement and cost $17M. When optimized, this same hardware will be able to attain Sequoia's stretch bandwidth goal. In conjunction with this effort, LC released a Linux port of ZFS to the public, which has been well received by the Linux community.

Progress in 2010

The cost and components of a Lustre file system can be greatly reduced by improving the method by which it uses back-end disk resources for its object store. Random disk accesses are slower than sequential accesses because of head seek latency ("seek time" is the length of time it takes a hard drive's read/write head to find the physical location of a piece of data on the disk, and "latency" is the average time for the sector being accessed to rotate into position under a head, after a completed seek). Unfortunately, the algorithms used by Lustre's back-end file system, EXT4, tend to result in random I/O to the disks when processing a production workload. To improve system performance and achieve other advances in data integrity, the LC–Oracle team replaced the back-end file system with the revolutionary ZFS, which was originally designed by Sun Microsystems, Inc. ZFS is a "copy-on-write" file system that effectively converts the concurrent, random writes arriving at an Object Storage Target (OST) to a stream of higher performing sequential writes. It also includes advanced caching technology that utilizes solid-state disks to improve read performance.

Zetabyte File System (ZFS) administrator Jim Silva works on one of the 10 disk enclosures in LC's prototype ZFS-based Lustre file system. The prototype has a capacity of 896 terabytes and delivered 8GB/s to Lustre clients in a 2010 demonstration. When this system is optimized, it is expected to deliver 16GB/s. The object storage portion of the prototype, comprised of 10 disk enclosures and eight servers, will be replicated 62 times to build the Sequoia file system.

LC championed the idea of combining ZFS and Lustre before Sun Microsystems acquired Cluster File Systems, Inc. (CFS) in 2007. For several years, LC has been leading the work of porting ZFS to Linux, while the Sun Microsystems team has focused on refactoring Lustre's back-end file system interface to allow customers to choose either ZFS or EXT4. This collaboration endured Oracle's acquisition of Sun Microsystems in 2010 and resulted in a public release of the ZFS port for the Linux community. LC demonstrated the integrated product on a prototype system in 2010.

LC created the Solaris Portability Layer (SPL) to port ZFS to Linux. SPL emulates Solaris kernel programming interfaces in terms of the facilities available in the Linux kernel. This architecture minimizes the changes needed to the ZFS source code and allows SPL to be separately tested with SPLAT (Solaris Porting LAyer Tests) and debugged. SPL is still undergoing optimization but is functionally complete, allowing ZFS to operate on Linux.

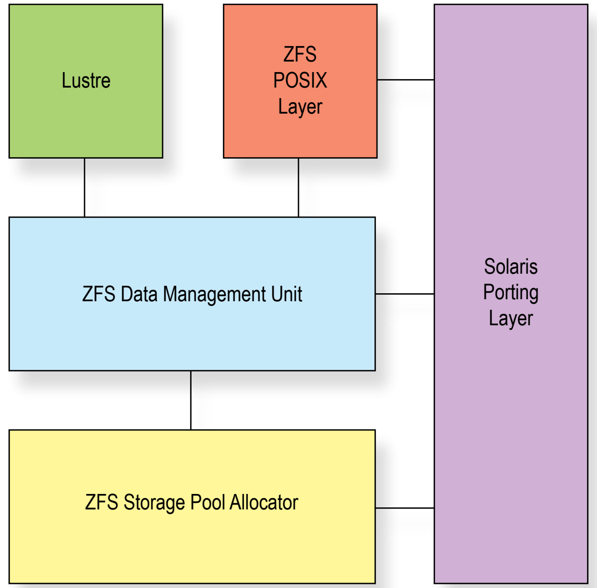

Lustre interfaces with the Zetabyte File System (ZFS) Data Management Unit (DMU), which provides a transactional object interface. DMU interfaces with the ZFS Storage Pool Allocator, which is akin to a logical volume manager. The Solaris Porting Layer provides Solaris kernel interfaces to allow ZFS layers to function in the Linux kernel. The ZFS POSIX Layer (ZPL), which is required for ZFS to function as a standalone file system on Linux, is not needed for Lustre, but members of the Linux community are working to port ZPL to Linux.

ZFS is implemented in layers. The Data Management Unit (DMU) is a transactional object interface that maps well to Lustre's OST model. Because Lustre interfaces directly with DMU, LC did not need to port the ZFS POSIX Layer (ZPL) to Linux. As a result, LC was able to focus on porting the DMU and Storage Pool Allocator layers. These layers became fully functional on Linux in 2010.

ZFS was designed to interface directly with inexpensive commodity disks. Its design for end-to-end data integrity checking, an integrated storage pool, and a unique RAIDZ algorithm allow it to manage redundancy and failures more effectively than an external hardware RAID appliance. For this reason, each Lustre server is directly attached to 140 spinning disks in the prototype. This server configuration presented some Linux scaling problems, as well as challenges to managing disk enclosures and drive faults directly from Linux, functions that are typically handled by a RAID appliance. LC solved several scaling issues and implemented basic enclosure management and a fault management framework. Partnerships for refining and extending this framework for a particular manufacturer's disks and enclosures are being sought via a Sequoia I/O research and development request for proposals.

The LC ZFS port generated significant interest within the Linux community when it was publicly released. In addition, community members volunteered to complete the ZPL port to Linux and to provide testing and feedback on various kinds of Linux hardware and distributions. These efforts strengthened the ZFS port while allowing LC to focus on its mission of building a viable file system for Sequoia. The strong response from the community, coupled with Sun Microsystems and CFS's release of their products under open source licenses, underscores the value of free software and open community development in the high performance computing system software arena.

In September 2010, LC's prototype system was installed with the latest Lustre and ZFS development code, attached via InfiniBand to the Sierra cluster, and exercised with the Interleaved or Random (IOR) parallel I/O benchmark. Using 32 compute nodes, one task per node, one file per task, and a 1 megabyte transfer size, IOR reached a write rate of 7.9GB/s. This trial run was the first end-to-end test of the new software on the prototype hardware. Optimization of the system is expected to increase the write rate to approximately 16GB/s. The prototype hardware has a peak capability of 25GB/s, limited by its eight quad data rate InfiniBand connections.

COMPUTATION DIRECTORATE ANNUAL REPORT 2010

RETURN TO TOC

OpenSFS Promotes File System Advancement and Community Participation

CONTACT
MARK SEAGER
SEAGER1@LLNL.GOV

An essential component of modern simulation environments is a global, parallel, and scalable file system that provides access to all of its computing resources, performs I/O operations simultaneously, and expands and adjusts as more I/O resources are brought online. To meet these requirements, the Laboratory's Advanced Simulation and Computing and Multiprogrammatic and Institutional Computing programs have made significant investments in developing, deploying, and supporting the Lustre scalable file system over the last 10 years.

Lustre is an open source development project that was initially funded by the Accelerated Strategic Computing Initiative PathForward program in the late 1990s and managed by LLNL. The initial development team included Hewlett-Packard, Intel, and Cluster File Systems, Inc. The Lustre technology was later acquired by Sun Microsystems, Inc., and then transferred to Oracle in 2010. In April 2010, Oracle announced that future development and support of Lustre would focus on Oracle products and Solaris operating systems. Consequently, LLNL embarked on a new effort to organize Lustre on Linux for the HPC community and create a 501(c)(6) California Mutual Benefit Corporation for providing ongoing support and development of Lustre. The organization is called Open Scalable File Systems, Inc. (OpenSFS). OpenSFS is a technical organization focused on high-end, open source file system technologies. The goals are to provide a forum for collaboration among entities deploying file systems on leading HPC systems, to communicate future requirements to the Lustre file system developers, and to support a release of the Lustre file system designed to the needs of the users.

Progress in 2010

Over the last 18 months, LLNL joined forces with Oak Ridge National Laboratory, Cray, Inc., and DataDirect Networks to establish OpenSFS. The lengthy organizational development process focused on building consensus within the HPC community for the OpenSFS business model, technical approach, and management structure. Attorneys then refined key documents to meet the legal and intellectual property objectives of the founding organizations and broader HPC community.

As a result of the outstanding responsiveness of LLNL's Procurement Department to prepare membership agreements and navigate the complex administrative processes, the Laboratory was the first entity to join OpenSFS. The other founding organizations became members approximately six weeks later. OpenSFS officially announced its formation on October 19, 2010, and the first community meeting was held at the 2010 Supercomputing Conference in New Orleans on November 16, 2010. Since that meeting, additional organizations have announced their intent to join OpenSFS.

The OpenSFS team formed four working groups in the following areas to begin the technical work of the user community: 1) technical, 2) release planning, 3) support, and 4) communications. Respectively, these groups are responsible for advancing Lustre development; coordinating development activities into a specific release schedule and performing testing and validation of the release; handling day-to-day software support defect tracking, resolution, and bug-fix testing; and maintaining the OpenSFS.org Web site, development tools, mailing lists, and communications strategies. The working groups must effectively collaborate with one another to successfully implement their organizational objectives. Other working groups will be formed as necessary.

LLNL's work to develop, deploy, and support Lustre on Linux for and in collaboration with the HPC community complements Oracle's commercial Lustre business plans. The ultimate goal is to provide a sustainable Lustre development and support environment that will meet the Laboratory's demanding programmatic requirements.

Open Scalable File Systems, Inc.

LLNL was the first entity to join OpenSFS, an organization that aims to develop, deploy, and support the Lustre scalable file system on Linux for the high performance computing community.

Linux Cluster Contributes to the Understanding of the Big Bang

CONTACT
JEFF LONG
LONG6@LLNL.GOV

Researchers recreated a miniature Big Bang on the Large Hadron Collider (LHC) at CERN in Switzerland in November 2010. Data from this landmark experiment is being analyzed using LC's Green Linux Compute Cluster (GLCC) and other systems on a worldwide grid dedicated to the ALICE (A Large Ion Collider Experiment) detector—one of four major detectors on LHC.

GLCC represents LC's first foray into grid computing. The cluster runs heavy-ion collision simulations and processes experimental data from ALICE. LLNL has teamed with Lawrence Berkeley National Laboratory (LBNL) and the National Energy Research Scientific Computing Center to create a major U.S. site on the worldwide ALICE grid, providing the primary computing and storage resources for more than 1,000 collaborators in North and South America.

Progress in 2010

Although fairly diminutive by LC standards, GLCC is one of the most productive systems on the ALICE grid. At any given time, GLCC is running up to 1,000 simultaneous simulation and analysis compute jobs and quickly producing results for which scientists around the globe are eagerly awaiting.

The GLCC compute component contains half of a scalable unit of the same hardware as LC's new 261-teraFLOP/s unclassified Sierra cluster. It has 48 gigabytes (GB) of memory per node and features the Westmere 2.8-gigahertz central processing unit. The storage component consists of SuperMicro hardware and totals 650 terabytes of usable storage.

The cluster represents two firsts for LC: the first attempt at grid computing at an international level and the first time that DOE's Office of Science has directly invested in LC-managed hardware. The project was spearheaded at LLNL by the Physical and Life Sciences Directorate, who secured funding through DOE's Office of Nuclear Physics within the Office of Science.

Development of the GLCC system architecture and other hardware requirements resulted from collaborations with ALICE researchers at LLNL, LBNL, and CERN. In addition to procuring, installing, and integrating the hardware, LC manages the system's day-to-day hardware operations and the grid software. GLCC resources are made available to the Open Science Grid when ALICE demands are low, providing a valuable computational resource for other researchers studying high-energy nuclear physics around the U.S.

The ability to site a grid computing system at LLNL was enabled by the existence of the Green Collaboration Environment, a special unclassified, unrestricted 10GB/s network that was created by LC to facilitate the sharing of data between LLNL scientists and external collaborators. The GLCC system and the Green Data Oasis, a 620-terabyte data storage system on the unrestricted network, are the two principal components of the Green Collaboration Environment.

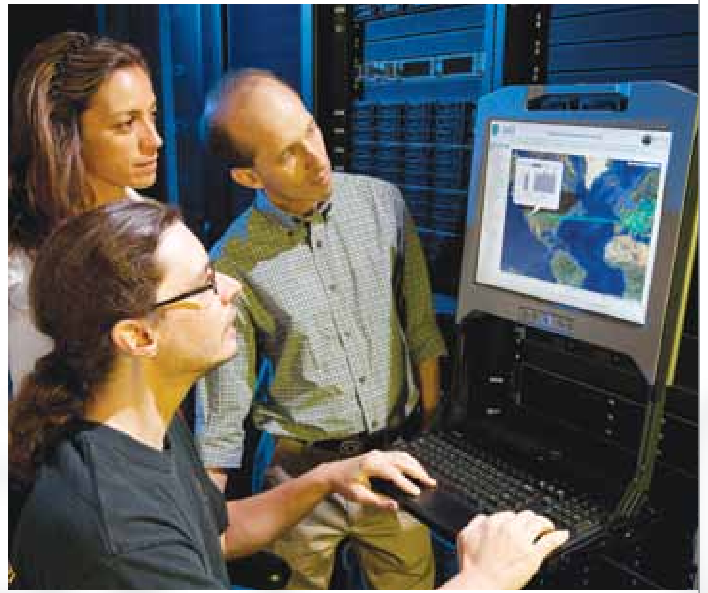

Collaborators monitor the use of the Green Linux Compute Cluster on the ALICE grid.

Computer Facility Earns National Recognition for Energy Conservation

CONTACT
ANNA MARIA BAILEY
BAILEY31@LLNL.GOV

The Terascale Simulation Facility (TSF) received a Leadership in Energy and Environmental Design (LEED) gold-level certification under the U.S. Green Building Council rating system on December 24, 2009. LEED, an internationally recognized green building certification system, provides third-party verification that a building or community was designed and built using strategies aimed at improving energy savings and water efficiency, reducing carbon dioxide emissions, and other factors. Although LEED certification is typically sought for buildings under construction, the completed TSF scored 56 out of 57 points on its submission, earning it the rare gold-level award.

In a congratulatory communication sent to LLNL, Brigadier General Garrett Harencak, the National Nuclear Security Administration (NNSA) principal assistant deputy administrator for Military Application, said, "This is truly a noteworthy achievement for NNSA and LLNL that symbolizes our commitment to transforming the Cold War-era nuclear weapons complex into a modern, efficient nuclear security enterprise."

TSF, which was completed in 2004, is a 253,000-square-foot building that houses several of the world's fastest supercomputers, including Dawn (BlueGene/P) and BlueGene/L. These Advanced Simulation and Computing systems are largely dedicated to stockpile stewardship.

Progress in 2010

The LEED certification process took approximately one year to complete. TSF was awarded points on its submission for many effective energy conservation efforts, including managing the airflow in the two 24,000-square-foot computer rooms, raising the air temperature in the computer rooms, raising the water temperature in the cooling system, and implementing an automated electrical usage system to reduce off-hours power consumption in the office tower. These individual efforts are part of the Computation Directorate's ongoing "Megawatts to PetaFLOP/s" program, which is designed to conserve energy in TSF by adjusting temperature and managing airflow.

The TSF project team also completed a 15-megawatt (MW) power upgrade of the east and west computer rooms, enabling the facility to achieve 30 MWs of computational capacity. The increased power distribution was achieved through the installation of 10, 1,500-kilovoltampere transformers; 20 switchboards; and four sectionalizing switches, as well as an associated duct bank, feeders, manholes, conduit, and elevated structural frame for conduit support.

The overall power upgrade was completed on schedule and under budget because of excellent project planning and execution.

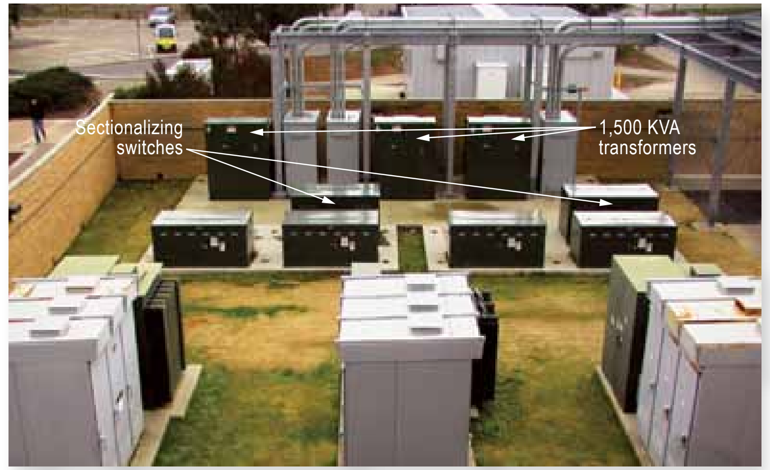

As part of the Terascale Simulation Facility (TSF) power upgrade, new 1,500-kilovoltampere (KVA) transformers were installed along with feeders and associated sectionalizing switches in the TSF west switchyard.

The project team accelerated material procurements for the upgrade by purchasing nearly 30 miles of copper wire when copper prices were favorable, saving an estimated $250K. The team also worked with the Department of Energy's Livermore Site Office to secure early approval of the Phase II addition to the project—construction on the east side of the building. As a result, construction and procurement activities began in FY09 rather than in FY10, which helped mitigate the risk of delays to the project.

This TSF upgrade marked the successful completion of a Laboratory Performance Evaluation Plan target that indirectly supports the 2011 acquisition of the Sequoia computer. In addition, these electrical projects will help extend LLNL's leadership in high performance computing and scientific simulation by providing the power capacity to site multiple world-class computers simultaneously at TSF.

Software Applications Help Build a National Advantage

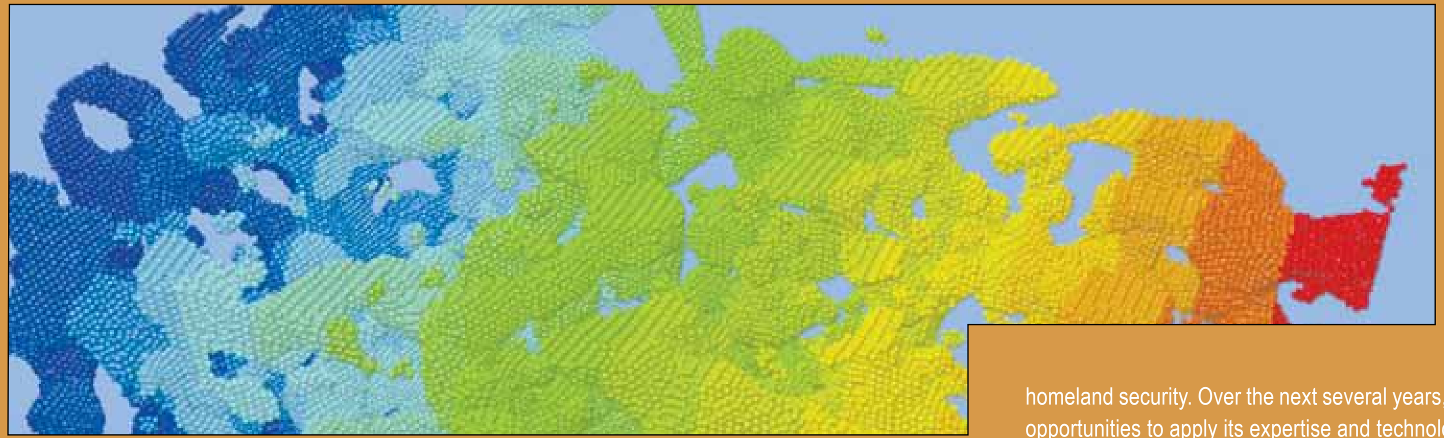

JOHN GROSH
GROSH1@LLNL.GOV

The ability to acquire and analyze data, make decisions, and take action faster than the competition or adversary drives economic, organizational, and military advantage. Examples abound in all disciplines: basic research is measured by first discoveries; military supremacy is acquired by superior situational awareness; market share in business is accomplished by making better decisions despite uncertainty. The Computation Directorate at Lawrence Livermore National Laboratory (LLNL) develops software applications that enable a national advantage by building tools that provide insight into physical and system phenomena, designing systems with faster time-to-solution, and analyzing system tradeoffs in great fidelity. The range of Computation's software applications spans from bioinformatics to real-time data acquisition and control to multi-physics simulations. The key to Computation's success is a foundation of software engineering, computer science, computational mathematics, and systems knowledge.

Historically, Computation's applications development activities have focused on supporting the Laboratory's federal government sponsors. LLNL's forte has always been developing multi-physics simulation codes that scale on massively parallel high performance computing systems. The article on performance studies on KULL, where computer scientists scaled simulations of radiation flow to more than 30,000 processors, is an example of this enduring capability (see page 34). In 2010, Computation computer scientists also developed knowledge management systems that support missions in intelligence, environmental protection, defense, and homeland security. Over the next several years, the directorate will more actively seek opportunities to apply its expertise and technologies to support industrial applications of high performance computing. This endeavor will involve developing partnerships with independent software vendors, manufacturers, and industry; adapting tools and applying expertise developed for nuclear weapons simulation; and addressing other significant problems in the private sector.

The challenges ahead demand increased insight in multiple domains. LLNL is aggressively building tools to model and analyze energy systems, which requires an adaptation of libraries and tools to new classes of application problems. Two examples include the analysis of wind turbine farms and design tools for retrofitting buildings for better energy efficiency. For these cases, Computation is developing applications software to understand fluid flow and heat transfer using Overture, a toolkit for solving partial differential equations on overlapping grids. These efforts involve collaborations with industry (e.g., United Technologies Corporation and Siemens Global) and leverage the Department of Energy Office of Science programs in computational science research. In addition, Computation is applying its expertise to parallel industrial codes for the future electric grid. The PLEXOS article in this section discusses the work of Computation computer scientists with Energy Exemplar and IBM to broaden insight into renewable power generation (see page 32).

Computation is building synergistic relationships that enable the interchange of software technologies, invigorate the directorate's existing core programs, and build new capabilities for the benefit of American economic competitiveness.

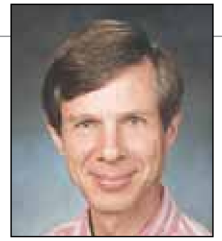

Cryogenic Target System Supports First Integrated Ignition Experiment

CONTACTS

OLIVER EDWARDS
ODE@LLNL.GOV
MIKHAIL FEDOROV
FEDOROV2@LLNL.GOV

The National Ignition Facility's (NIF) Cryogenic Target System (CTS) is a hardware–software component of the Integrated Computer Control System (ICCS). It provides precise cryogenic temperature control, x-ray imaging of the target along three axes, and coordinated process control and image analysis. CTS works in conjunction with NIF Shot Data Systems and has supported 61 NIF shots with cryogenic targets and the fabrication of nearly 50 target fuel ice layers. These efforts led to the successful first integrated ignition experiment in September 2010.

To achieve ignition, NIF cryogenic targets must contain an extremely uniform deuterium–tritium (DT) fuel ice layer. Attaining this precision has been a long-standing challenge for the Inertial Confinement Fusion program. CTS helps solve the issue by crystallizing layers from a DT seed and executing thermal protocols with better than 1-millikelvin (mK) precision. The CTS x-ray imaging system monitors the growth of the layer along three axes. Analysis algorithms enhance these images and evaluate the metrics of the 75-micrometer (µm)-thick layer against the 1-µm roughness specification. CTS also provides a unique environment that allows target physicists to directly interact with the layering process. Once the layered target is formed, CTS collaborates with NIF shot systems in a carefully orchestrated shot sequence. It then stores data in the long-term NIF archive for further visualization and analysis.

Progress in 2010

On September 29, 2010, NIF conducted its first integrated ignition experiment, achieving a National Ignition Campaign (NIC) Level-1 Milestone. CTS was a critical component in the success of this shot. CTS controls help create, measure, and maintain a layered ignition-capable target. The system also supervises more than 100 control points ranging from simple thermocouple sensors to high-performance x-ray imaging cameras.

CTS is based on ICCS frameworks and takes full advantage of ICCS's scalable, distributed Common Object Request Broker Architecture (CORBA). CTS leverages the cross-platform, cross-language interoperability of ICCS to integrate instrumentation hardware from dozens of vendors into one unified system. The flexibility inherent in the ICCS distributed architecture allowed optimal placement of computing resources in many locations such as the NIF control room, server rooms, instrumentation racks, and adjacent to the cryogenic equipment.

CTS utilizes a wide range of software technologies. Its front-end processors—which interface to the hardware and serve as adapters to supervisory-level controls—are coded in Ada and Java and run on real-time VxWorks, Solaris, Windows, and Linux operating systems. Middle-layer supervisory objects, including graphical user interfaces, process controls, and monitoring systems, are coded in Java on Solaris and Windows. The layering software is written in MATLAB and runs on Linux. The shot control system is coded in Ada and Java.

When preparing for and executing an ignition shot, CTS transitions between a number of states, such as cool down, layering, and post-shot warm-up. As it changes from one state to another, CTS configures temperature control and imaging engines with dozens of parameters for each state. The cryogenic temperature control engine provides precise (better than 1-mK accuracy) control of the target, fuel reservoir, shroud, and various thermal shields of the Ignition Target Inserter Cryostat (ITIC). (ITIC is attached to the end of the target positioner and cools the target and the DT fuel mixture to meet temperature and uniformity requirements.) The imaging engine acquires sequences of x-ray images along three axes and synchronizes imaging with the stable periods created by cycling off the vibrating cryocooler.

Traditionally, NIF controls are driven by discrete events centered around the NIF laser shot. However, cryogenic target layering is a continuous real-time control process with layer formation lasting 24 to 72 hours. During this time, sensor measurements and images of the capsule are continually acquired and analyzed, and outputs are updated. To support this work, CTS extends ICCS capabilities with tools suitable for continuous process control tasks, such as improved cryo-activity automation, periodic archiving, and a "stripchart" continuous data-monitoring application.

A unique challenge for CTS is accommodating the continuous flow of targets through NIF. A target's layering response—called a layering "recipe"—is measured and recorded in an off-line Ignition Target Proofing Station (ITPS) prior to installation in NIF. ITPS runs a version of CTS controls to ensure that the layering recipe translates with fidelity from ITPS to NIF.

CTS must accommodate various target types. It handles the variations with a data-driven design. By changing the input data, the system can be quickly and easily reconfigured, allowing operators to adjust the system between shots or even during a shot, and thus handle unexpected events like sensor failures without stopping the shot or restarting the control system. Maps of target control points are stored in an Oracle relational database management system. In the future, each target's map will be automatically loaded from NIF Enterprise Data Systems.

CTS is required to support layering algorithm research and to provide interactive layering controls capabilities. The system provides these functions through a unique interpretive environment, built on MATLAB, that allows physicists to directly interact with and modify the layering process. The CTS toolbox for MATLAB—based on CORBA interoperability—provides an abstracted set of objects to the MATLAB environment, enabling MATLAB applications to control and monitor CTS cryogenic temperature, x-ray imaging, visualization, and data systems. These virtual controls match the layering domain and isolate MATLAB codes from direct interaction with the control hardware.

In 2010, CTS layering and cryogenic controls were developed and deployed, successfully supporting dozens of cryogenic shots and culminating in the first integrated ignition experiment. In 2011, Computation teams will deploy software to automate cryogenic activities, thereby decreasing the operator workload and improving reliability. By integrating CTS with the NIF Computing Information Systems business applications Glovia and Locos, much of the configuration data will be automatically, rather than manually, loaded for each experiment. In addition, as layering becomes better understood, layering process controls will be transferred from the MATLAB environment to an ICCS graphical sequencer. Computation teams will also build post-shot reports and extend post-shot visualization capabilities.

The Cryogenic Target System provides precise cryogenic temperature control of the ignition target, its protective shrouds, and the supporting Ignition Target Inserter Cryostat (not pictured). This target was mounted in NIF for the first integrated ignition experiment.

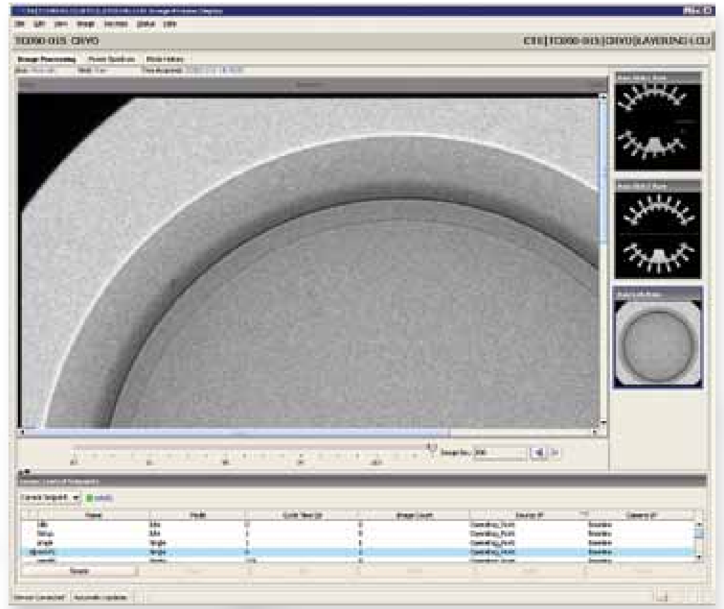

The Cryogenic Target System's layering graphical user interface provides real-time feedback on the status of the DT ice layer in the target. This image shows the high-quality DT ice layer used in the first integrated ignition experiment.

RETURN TO TOC

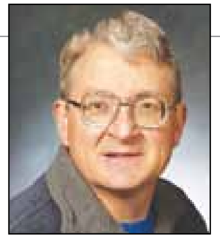

Bioinformatics Team Advances Pathogen Detection Research

CONTACT
TOM SLEZAK
SLEZAK1@LLNL.GOV

Bioinformatics research at LLNL began with work for the Human Genome Program in 1987. At that time, LLNL focused on building practical infrastructure and tools to store, analyze, and display information for human chromosome 19. Almost a decade later, LLNL helped establish the Department of Energy's informatics infrastructure for the Joint Genome Institute in Walnut Creek, California. Through this experience, the LLNL bioinformatics team developed an extensive, whole-genome mindset that has proved vital for overcoming bio-related challenges of this century. In 2000, LLNL was part of a team that built a wide-area biosecurity system for the 2002 Winter Olympic games in Salt Lake City, Utah. The LLNL team created the first automated pathogen DNA signature detection system, utilizing novel whole-genome comparison techniques. This work put LLNL in a leadership position for developing pathogen detection assays, which have been used since 2003 in BioWatch, a nationwide early-warning system that detects trace amounts of specific microorganisms in the air.

In recent years, LLNL has pioneered using high-density microarrays to produce high-resolution assays for biodefense applications, such as pathogen detection and forensics, and for public health and safety projects. The Lawrence Livermore Microbial Detection Array (LLMDA) confirmed previously unknown, and fortunately benign, viral contamination in a vaccine given to babies. Other microarrays accurately defined the strain-resolution of pathogens for the Department of Homeland Security's (DHS's) Microbial Forensics center, and sophisticated analysis software was developed and delivered to the DHS customer.

Other accomplishments this year include performing ultra-deep DNA sequencing on rabies viruses isolated from an apparent species jump from skunks to foxes. Assays for 27 bacterial causes of sepsis were developed for a company performing clinical trials in Australia. In addition, the bioinformatics team is currently building a Microbial Forensic "encyclopedia" database for DHS, as well as a protein reference infrastructure for a Department of Defense medical countermeasures program.

Progress in 2010

LLMDA is a glass microarray that contains 388,000 DNA fragments (probes) that allow LLNL scientists to identify more than 900 sequenced bacteria and more than 2,000 sequenced viruses. In early 2010, a collaborator at the Blood Systems Research Institute in San Francisco, California, asked the Livermore team to verify the presence of a contaminant that had been detected while sample sequencing eight vaccines typically given to small children. The team used LLMDA to confirm the presence of a pig virus that had gone undetected through two years of clinical testing and more than four years of worldwide vaccine usage.

The U.S. Food and Drug Administration (FDA) is examining how new technologies, including arrays, might best be used to improve product safety. LLNL's press exposure from helping with the vaccine issue led to discussions with several large-scale food and product companies that are now considering whether LLMDA might improve their ability to detect potential pathogens before they affect consumers.

LLNL also built and tested several forensic microarrays for the DHS Microbial Forensics center, which relied on LLNL for operational testing while their labs awaited formal certification. These arrays demonstrated that precise strain comparisons could be made to already-analyzed strains in 24 hours for $500, compared to the 72 hours currently required for DNA sequencing for $5K. Using newly developed LLNL software, the Laboratory team efficiently analyzed hundreds of pathogen genomes to discover single-base DNA changes that can provide forensic discrimination. DHS also funded LLNL's participation in a new collaboration with Australia. This team will use forensic arrays to screen 100 pathogen samples and determine which appear most novel and worthy of DNA sequencing. The new analysis software also allows researchers to combine forensic data from arrays and from genomic sequencing.

In a separate project, LLNL began working with members of the California Department of Public Health to investigate a recent outbreak of rabies in Northern California foxes that typically affects only skunks. The team is examining whether genetic mutations are responsible for the virus jumping from one species to the other. As part of this effort, new computational tools are being developed for high-throughput sequencing to examine both inter-host and intra-host viral evolution within the outbreak, including the role of rare genetic mutants.

In late 2009, a start-up company requested LLNL's help in designing traditional DNA signatures for the 27 known bacterial causes of sepsis (blood infections). The Laboratory team used specialized, LLNL-developed software to complete the assay designs in 2010 and predict their suitability for multiplexing into a single test. The assays are undergoing extensive testing in western Australia. The sponsor plans to seek FDA certification of the assays after they enter approved clinical use overseas. A diagnostic such as this can provide rapid and precise identification of the cause of sepsis, enabling physicians to prescribe the right antibiotics at the onset of the disease, which will ultimately save lives.

In support of the Defense Threat Reduction Agency (DTRA), the LLNL bioinformatics team contributed to a major program aimed at improving medical countermeasures. The team performed systems studies highlighting gaps in pathogen knowledge, and they were instrumental in helping DTRA obtain pathogen-sequencing contracts with several major research groups, including the Center for Disease Control and the University of Texas Medical Branch at Galveston. In addition, LLNL is building a software infrastructure to model protein structures in a large-scale pipeline. This work will help researchers perform analyses that could aid in the discovery of potential countermeasure targets and in designing candidate small molecules.

The bioinformatics team was tasked to integrate relevant knowledge of pathogen genomes, signatures, genes, microarray probes, and other data into a database that will be used for the DHS Microbial Forensics center. An initial version of this system was made available to DHS and new features continue to be added. This system could potentially enable multi-agency sharing of diagnostic and forensic data among contributing partners.

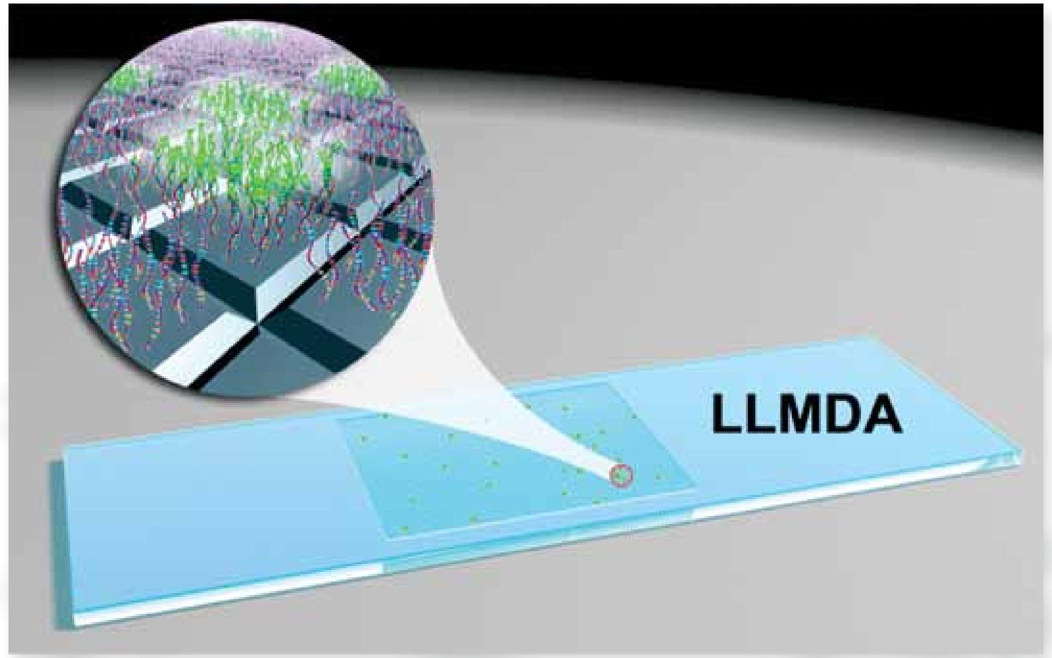

Microarrays contain discrete regions of DNA probes bound to a glass slide. When hybridization to a sample occurs, a fluorescent dye (green) is energized, indicating that a certain DNA signature has been observed.

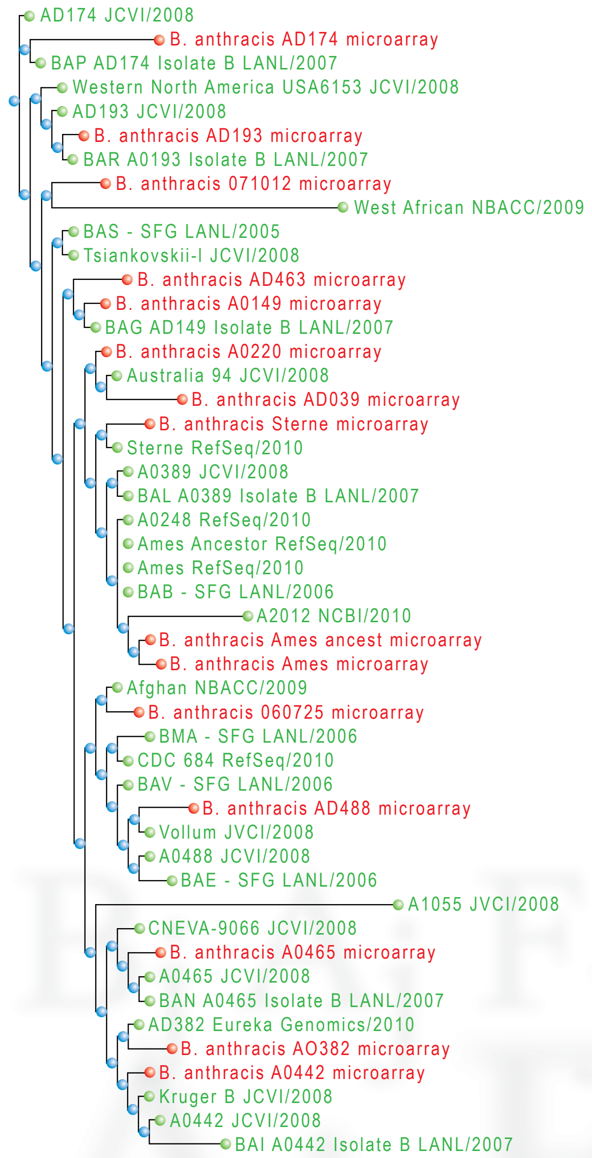

Forensic-resolution data from DNA sequencing (green) and LLNL microarrays (red) are shown in a phylogenetic tree, which permits rapid understanding of the significance of an outbreak or attack strain.

COMPUTATION DIRECTORATE ANNUAL REPORT 2010

RETURN TO TOC

Collaboration Leads to Advances in Power Grid Modeling

CONTACT
TODD GAMBLIN
TGAMBLIN@LLNL.GOV

As the California Public Utilities Commission (CPUC) aggressively works to integrate renewable energy sources into the power grid, understanding the effects of wind and solar power on the energy supply is becoming increasingly important. Computer models can provide an effective means of studying how these sources will influence the power market. However, even small power market simulations can require days to compute, with larger studies taking months or longer.

LLNL has partnered with Energy Exemplar and IBM to use high performance computing (HPC) resources for simulating the power grid. The collaboration has led to the development of a massively parallel implementation of Energy Exemplar's power market simulation software that has dramatically reduced the time required for large power market studies. As a result, studies that previously took months to run can be completed overnight. This preliminary work paves the way for more detailed grid models and a deeper understanding of clean energy sources for California.

Progress in 2010

Integrating renewable energy sources into the power supply is a complex problem because the amount of wind and solar power generated fluctuates based on the weather. To account for this variability in the future, conventional power plants must be added to the grid so that sufficient power continues to be generated even when wind and solar sources are unavailable. CPUC is currently conducting a study to determine how much additional conventional power generation will be needed in 2020 if 33% of the power grid is comprised of renewable sources.

Studies like the one undertaken by CPUC require sophisticated software to formulate models of the power grid. These models are used to generate mixed-integer programs for determining an optimal power generation schedule. Integer programs for renewable energy simulations are particularly large and difficult and require considerable computational time to solve.

In 2010, LLNL began a strategic partnership with Energy Exemplar and IBM to accelerate power grid simulations using HPC. LLNL worked with Energy Exemplar to develop a version of the company's PLEXOS power market simulation software that runs on Linux. The software was then coupled with IBM's CPLEX mixed-integer solver to build a massively parallel implementation of PLEXOS that can launch many concurrent simulations on LLNL clusters. This new capability has enabled LLNL to conduct previously intractable Monte Carlo studies.

In early November, the LLNL–private industry collaboration reached its first milestone: completing a parameter study of 1,000 concurrent PLEXOS simulations. The results provided new insights into ancillary service shortages and prices, which do not follow normal distributions and require very high sample sizes to estimate. Without the use of HPC resources, this type of analysis would have taken months, but with LLNL's Hyperion cluster, the study was completed overnight.

LLNL will continue to collaborate with Energy Exemplar and IBM to develop new power grid analysis techniques using HPC. The Laboratory is currently working toward larger scale simulation capabilities with PLEXOS, and it is beginning to investigate ways to extract more parallelism from IBM's mixed-integer solvers. This work has the potential to significantly improve understanding of renewable power generation and to better prepare the State of California for integrating renewable energy into its electrical grid.

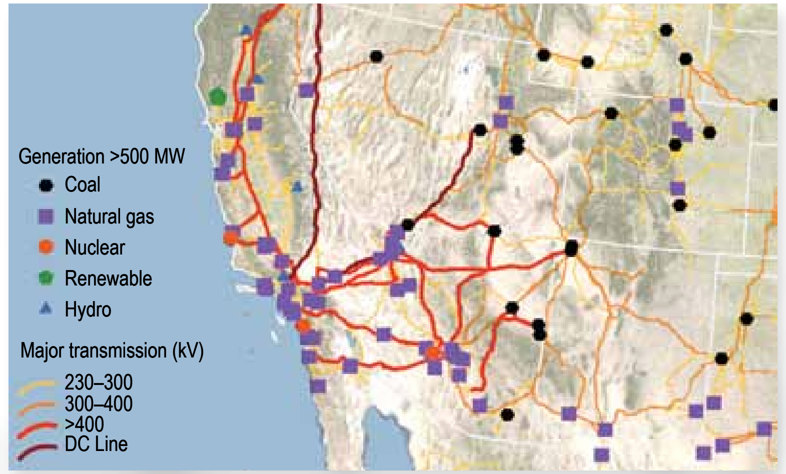

An energy map for the western region of the United States shows power plant locations and transmission lines for coal, natural gas, nuclear, hydroelectric, and renewable power plants.

RETURN TO TOC

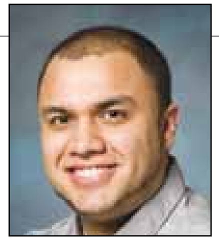

LLNL Improves Institutional Geographic Information Systems

CONTACTS
DEBORAH DENNISON
DENNISON3@LLNL.GOV
JOEL MARTINEZ
MARTINEZ248@LLNL.GOV

Geographic Information Systems (GIS) have become ubiquitous as evidenced by the widespread use of smartphones, Internet mapping services, navigation global positioning systems, and derivative applications. GIS maps provide an intricate network of information, unifying peoples' perspective of the world and enabling visually compelling views of diverse spatial data.

Several LLNL organizations recognized the importance of GIS and integrated the technologies into their business strategies by building geospatial services and applications. These advances provided the basis for new work-for-others projects and made previously cost-prohibitive activities possible. Recently, LLNL organizations have started to share their geospatial resources with one another, resulting in unanticipated benefits that have enabled LLNL to develop solutions for real-world, mission-related challenges.

Progress in 2010

Current GIS capabilities exist in LLNL's Global Security Principal Directorate through the Homeland Defense Operational Planning System (HOPS), National Atmospheric Release Advisory Center, and Joint Conflict and Tactical Simulation programs. The capabilities are also an integral part of the Environmental Restoration Division (ERD); and the Environment, Safety, and Health (ES&H); Facilities and Infrastructure; and Engineering directorates. (Computation researchers and scientists are embedded in each of these organizations.) However, most GIS organizational efforts are specifically tailored to meet individual programmatic needs.

The next step for advancing GIS at LLNL is developing a shared aggregation/distribution infrastructure that can serve as a hub for spatially enabled information. Many LLNL programs stand to benefit from an institutional GIS capability, including ERD's GIS Environmental Analysis and Visualization Tool (GAVT). GAVT accesses real-time programmatic data but is limited to static copies of institutional information such as LLNL base maps and topography. An institutional infrastructure would allow direct access to real-time external data sources, such as site facility information and satellite imagery.

Global Security's HOPS project has developed prototypes for many of the necessary components that comprise a GIS hub. HOPS enables sponsors, including the California National Guard, to manage risks associated with emergency response by providing environmental context. HOPS combines diverse and distributed information sources, such as populations and critical infrastructure, to synthesize a location-specific perspective of a given problem. This system transforms dynamic and static information into interactive, real-time, map-based applications and services to visualize risks and assets. Collectively, this system creates situational awareness from which users can coordinate common operational strategies.

HOPS is proving to be a transformative capability that will serve as the model for future LLNL institutional GIS standards and infrastructure. In a grassroots effort, ERD and ES&H are sharing spatial data with HOPS that augments existing emergency response data feeds. Future distribution facilities will be developed to share this information with several programs. Ultimately, these capabilities will be expanded to organizations outside of the Laboratory. After successfully demonstrating the GIS hub approach, a Computation team expects to implement a Laboratory-wide version based on the HOPS model and lessons learned from its use.

An institutional GIS retains data stewardship in the programs while capitalizing on the collective infrastructure. GIS is a natural mechanism for integrating specific programmatic data into a collaborative framework, allowing researchers to benefit from terabytes of information. As a result, programs will be better able to leverage these resources to innovate and deliver solutions to the nation's challenges.

The Data Integration and System Collaboration application provides rapid access to inter-organizational data feeds such as building layouts, chemical inventories, and facilities information. Clicking on an LLNL facility (shown in red) pops up a window with the available information for that building.

KULL Scaling Study on Dawn Expands Computational Capabilities

CONTACT
TYLER J. ALUMBAUGH
TALUMBAU@LLNL.GOV

Dawn, a 500-trillion-floating-point-operations-per-second (teraFLOP/s) BlueGene/P system, will lay the applications foundation for computing on Sequoia—the 20-qaudrillion-floating-operations-per-second (petaFLOP/s) system that will be deployed in 2012. The KULL application, a massively parallel, three-dimensional inertial confinement fusion simulation code, was ported to Dawn in 2010. This port enabled two important achievements in running KULL on advanced architectures: completing a scaling study in which more than 32,000 Message Passing Interface (MPI) tasks were performed, and successfully executing a large multi-physics simulation where full restart was achieved with 18,000 MPI tasks.

Progress in 2010

For the scaling study, the KULL team ran a simulation script that included only one physics package. The simulation, which involved radiation flowing through a "crooked pipe," tested the KULL infrastructure and the Radiation Diffusion physics package. In this "weak" scaling study, the zone-count-per-mesh domain remained approximately equal for every simulation.

The simulation ran for 50 time steps, and timing information was recorded for each cycle. The team performed the runs in symmetric multiprocessor (SMP) mode (one MPI process per node) and, fortunately, found that the code also ran acceptably in dual mode (two MPI processes per node). By running the entire scaling study in dual mode, the team avoided needing additional resource allocations. Initially, the results showed that the timing of the Radiation Diffusion physics package was dominated by a call to a parallel linear solver routine (from the *hypre* library) that demonstrated poor scaling. After consulting with the *hypre* developers, the KULL team adjusted the settings for the preconditioner routines, which significantly improved scaling.

The second study involved running KULL in a more complicated simulation utilizing larger processor counts than previously attempted. This simulation used the Lagrangian hydrodynamics, Arbitrary Lagrangian Eulerian remap, and a different radiation transport package using the method of discrete ordinates (or S_n). The mesh was created and partitioned into 18,000 domains. While attempting to run this problem using Dawn's dual-mode capability, the system's memory limit was exceeded shortly after beginning the simulation. Therefore, the KULL team ran the simulation using SMP mode, which performed 18,000 MPI tasks. The simulation ran for 24 hours and completed 83 time steps.

In total for the second study, three sets of complete restart files were generated (cycle 0, 30, and 60)—one per MPI task. Each restart file was between 422 and 500 megabytes, equaling approximately 8.3 terabytes of data per restart dump. The time to dump each complete restart set was 153.6 seconds, 494.2 seconds, and 519.4 seconds, respectively. These numbers represents an aggregate input/output (I/O) bandwidth of 55.33 gigabytes per second (GB/s), 17.20GB/s, and 16.36GB/s. As a percentage of peak I/O bandwidth, restart times were 81.3%, 25.3%, and 24.1% of the best performance deemed practical on Dawn. The cause of the widely varying I/O rates is still under investigation.

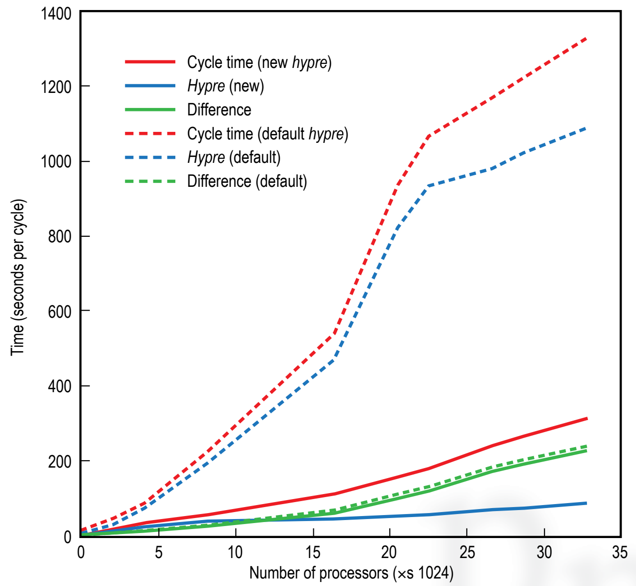

This plot shows the improved average cycle time versus processor counts on Dawn for the "crooked pipe" scaling study.

Information Technology Improvements Begin an Era of Technology Transformations

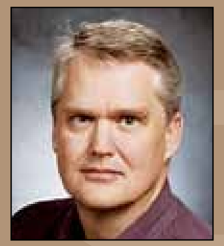

DOUGLAS EAST
EAST1@LLNL.GOV

Information technology (IT) should no longer be viewed simply as a set of devices and software tools used to conduct modern business but rather as a strategic mechanism that can dramatically increase the productivity of the workforce, as well as contribute to efforts to recruit and retain that workforce. Indeed, IT has embedded itself into the very fabric of modern life through numerous advances, particularly in communication technologies such as the Web, social media, cellular telephones, and wireless connectivity. In the past, technology introduced into the business environment seemed to drive the desires of what users wanted at home. Now, the ever-increasing dependence on technology in everyday life is driving users' expectations of what should be available at work. This new reality poses a significant challenge to Lawrence Livermore National Laboratory (LLNL), which operates under unique budget and security constraints. The federated model of IT support at the Laboratory also complicates the deployment of technologies.

Laboratory and IT managers' fiscal decisions must be made by striking a balance between maintaining an aged IT infrastructure, making investments to upgrade the infrastructure, and enabling new services. In addition, limited manpower means that the drive to deploy new technology is in competition with the necessity of maintaining services already in production. Because LLNL employees are critical contributors to many national security programs and missions, IT decisions must also be made in accordance with ever-present security concerns. The security landscape is further complicated by the Laboratory's vital unclassified academic and industrial collaborations—collaborations that would benefit from a more open security posture.

Despite these challenges, Computation employees significantly improved LLNL's IT services in 2010, as evidenced by the articles in this section. However, the Safeweb pilot (see page 38), the deployment of Microsoft Exchange (see page 41), and the new Blackberry service, though all important, have provided only a glimpse of the IT improvements that are necessary for the Laboratory to remain competitive in fulfilling its missions, meeting sponsors' needs, and recruiting and retaining an outstanding scientific workforce. In the future, enabling mobile computing, completing next generation networking, expanding the availability of wireless, and unifying communications (voice, video, and data) will be critical efforts aimed not merely at modernizing LLNL's IT capabilities but in transforming the way the Laboratory conducts business. In all of these efforts, Computation will play a major role through its divisions that support the Chief Information Officer (CIO) Program.

Infrastructure Upgrades Increase Workforce Mobility

CONTACT
ROBYNE TESLICH
TESLICH2@LLNL.GOV

Mobile computing is becoming the primary information access and application vehicle for consumers and businesses. Device form factors are continuing to evolve to offer a broader range of computing capability on highly mobile devices. In addition, application vendors are targeting mobile platforms as their primary devices, leaving the standard desktop as a secondary recipient. The popularity of the iPad and iPhone clearly shows that consumer applications are moving away from the desktop. As handheld computing devices become more powerful and as wireless connectivity becomes more pervasive, small form-factor devices will affect every aspect of our lives. In some cases, users are already relying on portable devices such as smartphones, Personal Digital Assistants, tablets, and laptops more than their primary desktop computers.

The ubiquitous use of mobile devices for information access in the consumer space is creating demand for the same level of information access and mobility in the workplace. To remain competitive, LLNL employees need secure access to information and applications from many form factors in many different locations (e.g., work, home, coffee shops, hotel networks). The Computation Directorate, in support of the CIO Program, continues to evolve LLNL's IT infrastructure and services to meet the demands of the emerging mobile trend.

Progress in 2010

Three major projects were completed in 2010 that provided the infrastructure necessary to more effectively mobilize the LLNL workforce: the Laboratory-wide migration to Microsoft Exchange as the e-mail and calendar system (see page 41); the service initiation and distribution of BlackBerry smartphone devices; and the risk assessment and creation of policy, standards, rules, and regulations that allow BlackBerry devices into many Limited Access areas at the Laboratory (see page 42).

In FY10, the CIO Program and Computation completed a massive multiyear project to convert LLNL's e-mail and calendaring system from Post Office Protocol/Meeting Maker to Microsoft Exchange. Exchange provided the unified communications solution necessary to deploy the first-ever LLNL-managed smartphone devices to employees. With more than 1,000 BlackBerrys now deployed, LLNL users are benefiting from the ability to access e-mail, calendars, task lists, and the Web from locations at the Laboratory and off site.

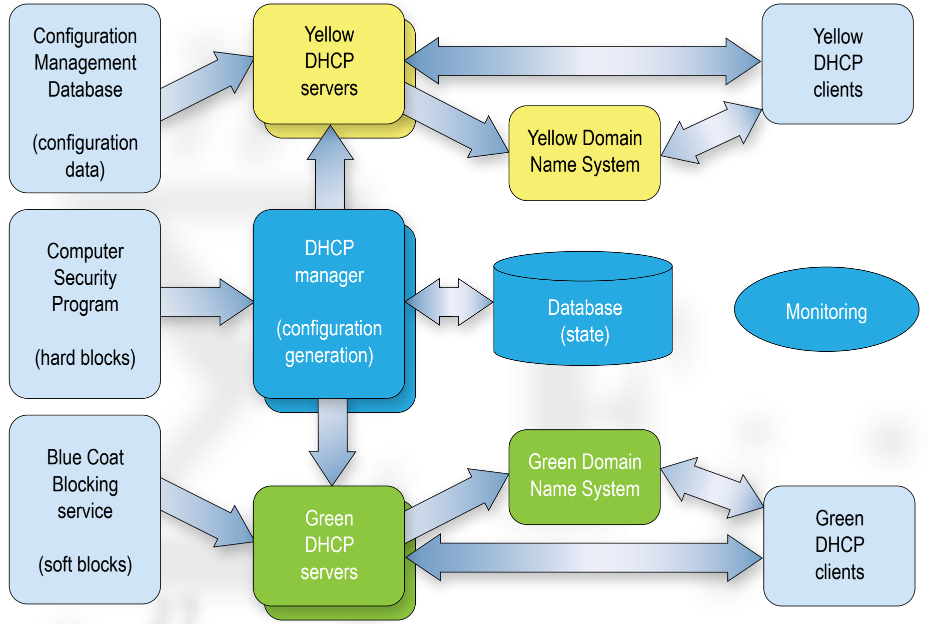

Dynamic Host Configuration Protocol (DHCP) service will allow mobile client devices to self-register with the network to obtain an Internet Protocol address.

Support for secure mobile computing requires a modern IT infrastructure. Unfortunately, technologies that these mobile platforms assume exist, such as Dynamic Host Configuration Protocol (DHCP), often do not exist in LLNL's IT environment. LLNL's strict security policies require additional conditions or restrictions on implementing secure access to network resources from mobile devices. Computation employees are working on several projects to address future mobile computing needs, including deploying DHCP, creating the LLNL Captive Portal for conference rooms, and expanding wireless networking to additional buildings.

To achieve true mobility, devices must be able to move from network to network, whether in an LLNL conference room, at a coffee shop or library Wi-Fi hotspot, or at home. In each location, the device must register itself with the network and obtain an Internet Protocol (IP) address using DHCP. Until recently, LLNL has refrained from implementing DHCP and has instead manually assigned and registered IP addresses to devices prior to allowing them to connect to the LLNL unclassified, restricted (yellow) network. In FY11, the project to deploy the DHCP service to LLNL networks will be completed. DHCP will enable devices to be configured to automatically register their IP in some Laboratory locations and facilitate easy access to LLNL networks while still complying with security operations. Once deployed, devices moving from one building to another will no longer need a manual reconfiguration of the IP address on the system.

Wireless access is available in many LLNL facilities today, and the capability will be expanded to other buildings if requested by the programs. While wireless deployment is growing, it is not yet available in all conference rooms nor can it be provided in Limited Access areas. Computation's Enterprise Network Solutions Group has developed a low-cost solution to provide secure wired access from a conference room to any employee or collaborator. The Captive Portal Network (CPN) service provisions users to the proper network based on their role and authorization. Users authenticate to the CPN, and a policy determines whether the user is provisioned access to either the unclassified, restricted (yellow) network; the unclassified, unrestricted (green) network; or the restricted (blue) network enclave.

The increasing mobility of the LLNL workforce combined with the convenience and availability of small form-factor devices are driving improved services and infrastructure in the LLNL IT environment. LLNL's planned path forward includes increasing wireless access, upgrading to the faster 802.11n standard of wireless network technology, expanding support for additional mobile platforms, and deploying technologies such as DHCP and secure yellow network conference room access for on-site mobile users. These improvements will significantly increase user productivity.

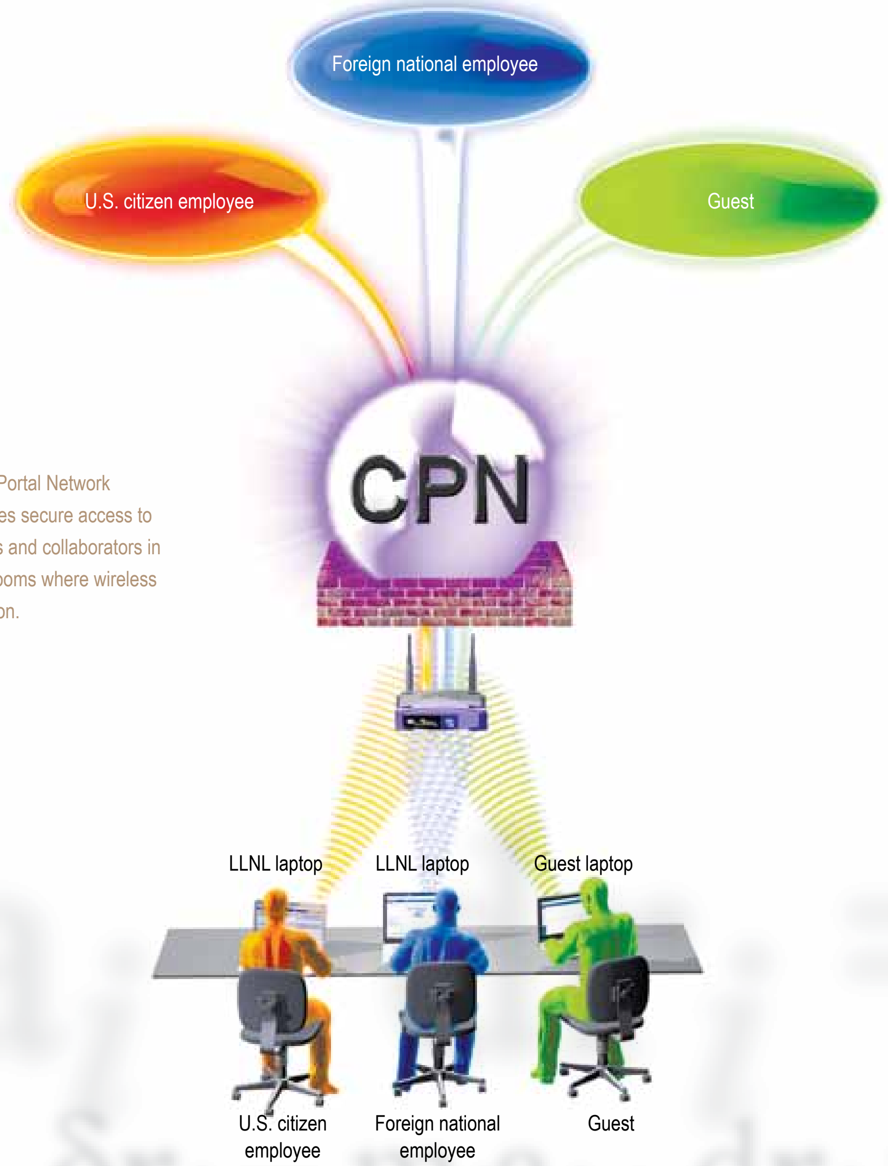

The Captive Portal Network (CPN) provides secure access to all employees and collaborators in conference rooms where wireless is not an option.

COMPUTATION DIRECTORATE ANNUAL REPORT 2010

Safeweb Improves Employee Productivity and Protects LLNL Networks

CONTACT
ROBIN GOLDSTONE
GOLDSTONE1@LLNL.GOV

In the summer of 2008, increased cyber threats prompted LLNL to impose significant restrictions on access to external social media Web sites, including Webmail services such as Gmail and social networking sites such as Facebook. These restrictions adversely affected employee work/life balance and prohibited employees from using social media for business-related collaboration.

A project called "Safeweb" was launched in early 2010 to assess the use of virtual desktop technology as a way to provide employees with access to social media sites while minimizing risk to LLNL's cyber security posture. Several technologies were evaluated, and a pilot was conducted with more than 1,200 LLNL users. The Safeweb solution has been overwhelmingly popular with pilot participants—so much so that Safeweb will be deployed as a production service in early 2011.

Progress in 2010

The Safeweb project began in February 2010. The project team established several basic requirements: the solution had to implement isolation to protect LLNL computers and networks; LLNL data could not be stored persistently in the Safeweb environment; the service had to be accessible by LLNL users from Windows, Macintosh, and Linux computers; the service had to be easy to use and perform "acceptably" as judged by pilot participants; and the cost needed to be justified by a positive impact on mission and cyber security posture.

The initial Safeweb concept was envisioned to utilize virtual desktop infrastructure (VDI) technology, where desktops are created as virtual machines on a centralized server. With VDI, users are issued a virtual desktop they access via a remote desktop client. Typically, a VDI solution provides users with their own personalized desktop, including access to stored data. In some cases the solution even remembers the last state of the desktop, allowing the user to connect to it from any location and pick up where they left off.

However, Safeweb adds a twist to the standard VDI model. A Safeweb virtual desktop is purposely destroyed when the user logs out, and a new, pristine desktop is created in its place from a master image. A pool of non-persistent desktops is available to a group of users, with the assumption that users access Safeweb occasionally throughout the day and thus are able to share the desktops and keep the solution cost down. A pilot was deemed necessary to validate this model and help determine an acceptable ratio of users to desktops.

The Safeweb project got off to a slow start when a lack of available IT resources precluded the implementation of a VDI-based solution. Instead it was suggested that the Safeweb concept could be implemented using simpler Windows Terminal Services (WTS) technology. A WTS-based solution was prototyped, and a pilot was launched in March 2010 with approximately 120 users. The pilot ran for three months, during which time a number of issues were encountered. WTS is a shared user environment, and as such, performance could not be acceptably controlled on a per-user basis. Additionally, because the solution was implemented inside the LLNL network, not all of the Laboratory's external Web site restrictions could be safely lifted. This issue limited the usefulness of the solution.

The WTS pilot thus provided two useful data points: a VDI-based solution was preferable, and an off-site implementation would provide optimal protection of LLNL assets. Armed with this information, the Safeweb project team began investigating how to implement Safeweb as a "cloud-hosted" virtual desktop service. Cloud computing refers to the ability to provision on-demand computing resources. Cloud resources can include servers, desktops, applications, or infrastructure. The on-demand nature means the customer simply specifies a desired capacity, and the provider is responsible for implementing and provisioning the resource. In the government sector, cloud computing often raises concerns about data protection. But because Safeweb, by design, does not store any LLNL data, it provided an ideal opportunity to explore the cloud computing model.

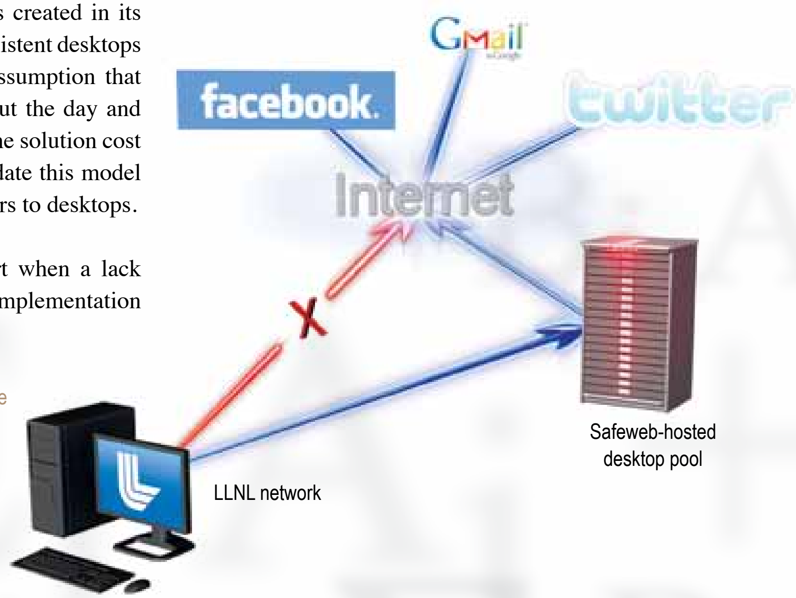

The Safeweb conceptual design illustrates how the service provides employees access to social media sites while minimizing risk to the Laboratory's cyber security posture.

The Safeweb project team spoke with several hosted desktop vendors and selected tuCloud.com as a vendor partner. Based on Safeweb requirements, tuCloud recommended the Kaviza VDI solution. Kaviza utilizes the Citrix HDX protocol, which is designed to optimize remote desktop performance across Wide Area Network (WAN) environments. Kaviza engineers worked directly with tuCloud and LLNL to implement a Safeweb prototype that was then evaluated by a small set of LLNL users. The Citrix client worked well on a variety of LLNL desktop platforms, and performance of the remote desktops, located at a data center in Southern California, was impressive.

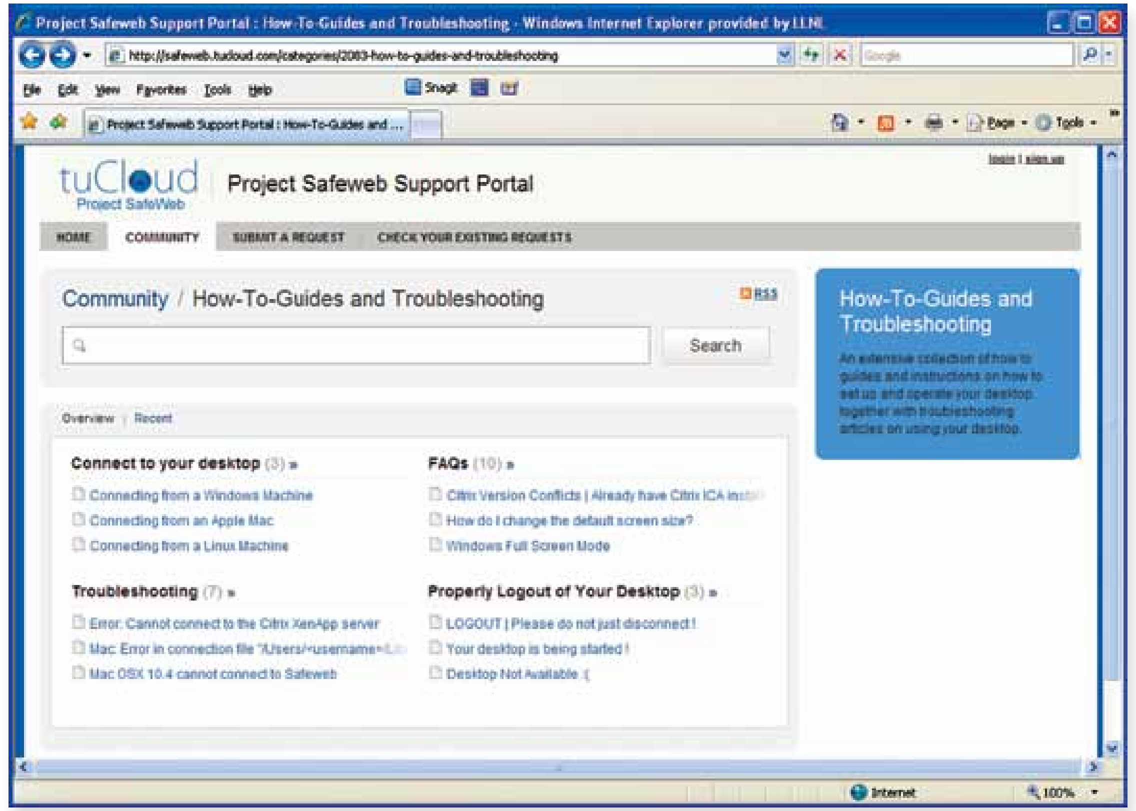

The Safeweb support portal provides users with how-to guides and troubleshooting scenarios.

Based on this successful prototype, LLNL signed a contract with tuCloud to implement a pilot with 250 desktops for 1,250 users for three months. User documentation was developed, the Citrix software was made available to LLNL users via the LANDesk software distribution portal, and a Web site was created to allow users to self-enroll in the pilot. The original Safeweb users were transitioned to the new service on August 30, 2010. Enrollment increased incrementally over the course of a month to ensure that pilot users received the assistance they required to start using Safeweb and to ensure that the service was performing and scaling as expected.

The Safeweb pilot has proven to be extremely popular—participant slots were filled shortly after the pilot was announced. In mid-October, a survey was conducted to assess how well Safeweb was meeting user expectations for functionality and performance. Responses were overwhelmingly positive. In fact, the most common criticism was that users would like to have remained logged in for a longer period of time. Many survey respondents cited legitimate business uses for Safeweb, while others pointed out the importance of staying connected to their personal lives outside the Laboratory. The survey feedback helped form the business case that is foundational to the funding request for ongoing operation of the Safeweb service. Safeweb will be scaled to 1,000 desktops and made available to all LLNL users in early 2011.

Next Generation Network Architecture Improves Reliability and Performance

CONTACT
LOUELLA PANAGA
PANAGA1@LLNL.GOV

The term Next Generation Network (NGN) connotes the end-state network architecture and implementation that has been developed for LLNL's unclassified networks. The goal is to provide a highly reliable and secure 10 gigabits per second (Gbps) converged backbone network capable of supporting data, voice, and video services. The NGN project consists of consolidating multiple physical networks, applying the latest technologies and best practices, and incorporating network designs that improve LLNL's security posture. NGN will enable a cost-effective delivery of next-generation services, such as Voice over Internet Protocol, Internet Protocol television, and mobility-related services; improve the overall reliability and performance of the network; and position LLNL to meet current and future network bandwidth requirements in support of programmatic missions.

Progress in 2010

Currently, LLNL's unclassified networks run on a 1Gbps backbone with dedicated physical infrastructure for each functional network segment: unclassified, restricted (yellow and blue); unclassified, unrestricted (green); secure, isolated "demilitarized zone" (DMZ); and wireless. NGN provides a starting point for LLNL to move away from the legacy infrastructure into a newly designed backbone that provides 10Gbps bandwidth from the distribution layer to the ESNet connection in a hierarchical network model. The core, distribution, and DMZ routers are configured with Cisco System, Inc.'s virtualization technology, Virtual Switching System (VSS), which allows two physical network switches to operate as a single logical entity. VSS reduces the operational complexity typically associated with high-availability configurations. Hot Standby Router Protocol, LLNL's current active/passive failover method, is replaced with an active/active configuration that increases available bandwidth and delivers stateful failover, which allows uninterrupted network connection to applications that rely on network state information, and deterministic sub-second failure recovery. The Open Shortest Path First protocol, a highly functional non-proprietary internal gateway protocol, replaces Cisco's Enhanced Interior Gateway Routing Protocol as the default routing protocol.

NGN implements a Cisco technology known as Virtual Routing and Forwarding (VRF), which permits multiple instances of routing tables to coexist in a router and operate simultaneously. VRF is a key architectural component that will allow the creation of multiple network security enclaves within the institutional unclassified network. NGN will replace the existing blue network firewall technology with a new, general-purpose "intra-enclave" firewall that permits instantiation of additional security enclaves. NGN also provides an overhaul of the LLNL Institutional Wireless Network service, adding support for the faster 802.11n standard and a new wired Captive Portal Network capability that extends employee and guest mobility to locations at the Laboratory where wireless cannot be deployed.

The NGN project team developed an integrated security architecture that satisfies LLNL security requirements while enhancing network efficiency and functionality. Their effort included designing a monitoring infrastructure that supports the ability to capture 10Gbps data flows and an improved perimeter firewall that consolidates the use of six Juniper NetScreen firewalls into two highly redundant Cisco ASA firewalls.

In 2010, the NGN project team completed the NGN architecture and design and installed the majority of the new network components. They also upgraded LLNL's main ESNet connection from 1Gbps to 10Gbps. In early 2011, the team will import security configurations from the current backbone to NGN and gradually migrate core services and building network segments to the new environment. NGN complements the CIO Program's initiatives to deliver advanced, innovative IT services to LLNL users.

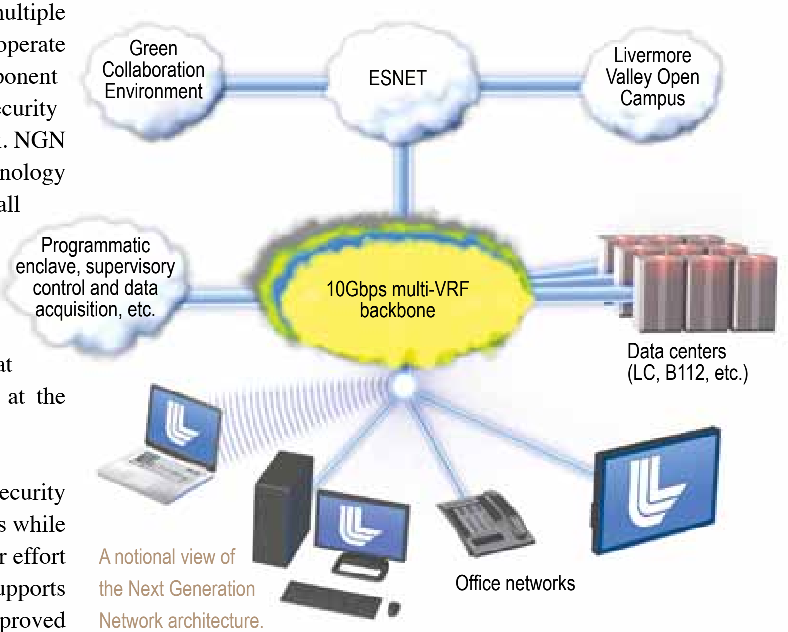

A notional view of the Next Generation Network architecture.

RETURN TO TOC

Exchange Migration Paves the Way for a Unified Communication Architecture

CONTACT
REA SIMPSON
SIMPSON6@LLNL.GOV

In 2010, Computation staff completed a multi-year project to migrate LLNL's e-mail and calendaring services from Post Office Protocol (POP) and Meeting Maker to Microsoft Exchange. The transition to Exchange, the market leader for enterprise-level e-mail and calendaring tools, positions the Laboratory to leverage myriad products from Microsoft and other third parties to provide a richer collaboration environment. The Exchange migration project was one of the largest IT projects ever undertaken by Computation's IT support organizations. The project replaced critical pieces of software that more than 6,400 LLNL desktop users access in their daily work.

Progress in 2010

The Exchange migration project consisted of two phases. During the first phase, which occurred over several months, users initiated an e-mail migration that transferred their accounts and server-side e-mails to the Exchange server. Computation staff developed the majority of the software used during the transfers, including a desktop tool that helped users convert the e-mail stored locally in Eudora to Outlook (PC users) or Entourage (Macintosh users).

The second phase of the migration was a simultaneous cutover of the calendaring service in August. For this phase, LLNL partnered with an external vendor, who provided the calendar migration tools. LLNL invested approximately 1,000 hours to testing and verifying the data migration process. During the actual 48-hour calendar cutover, approximately 325,000 calendar items were successfully converted from Meeting Maker to Exchange. Only 12 items required special attention.

The Exchange calendaring system works very differently from Meeting Maker. To help users cope with the transition, Computation staff developed detailed Web documentation comparing common tasks in Meeting Maker to common tasks in Exchange. Because calendaring functions are highly critical to their job responsibilities, LLNL administrative staff participated in hands-on Computation-led training sessions to prepare them for the changes.

LLNL is already benefitting from the migration to Exchange. Beyond the general efficiency gains from an integrated e-mail and calendar tool, Exchange provided the support necessary to deploy e-mail and calendar access on BlackBerry smartphones. In addition, Laboratory employees can now send and receive meeting invitations to and from people external to LLNL. These benefits have enabled more efficient meeting scheduling and management.

Exchange provides the foundation for a unified communication architecture. LLNL's instant messaging (IM) system, Office Communication Server, integrates with Exchange and provides new options to LLNL users. For instance, an e-mail message can be initiated from the IM client, and an IM conversation can be initiated from an e-mail message. Also, IM conversations can be archived in users' Exchange mailboxes. Computation is pursuing several projects that will provide additional unified communication features, including integrating the IM client with the LLNL IP-based phone system. Once complete, the IM presence status (the colored icon denoting "Available," "Busy," "Away," etc.) will include a status of "On the phone," and voicemail messages will be delivered directly to a user's Exchange mailbox.

In 2011, LLNL plans to upgrade from Exchange Server 2007 to Exchange Server 2010, which will provide an enhanced Web client for Macintosh and Linux users. In addition, Outlook for Mac will replace Entourage to provide an improved experience for Macintosh users.

The complex, multifaceted migration to Microsoft Exchange required careful planning, coordination, communication, and implementation. Approximately 45 multidisciplinary people from across the Laboratory had major roles in this project.

COMPUTATION DIRECTORATE ANNUAL REPORT 2010

Cyber Security Improvements Help Protect Institutional Assets

CONTACT
EDWARD MATSCHE
MATSCHE1@LLNL.GOV

LLNL must maintain a secure computing environment to accomplish its mission. Therefore, supporting LLNL's cyber security organization is one of Computation's most important responsibilities. Directorate staff devote significant effort and have achieved considerable success in providing timely, effective, and cutting-edge computer security support to the Laboratory. In 2010, Computation personnel improved LLNL's cyber defenses by installing new antivirus software; coordinating risk assessment efforts that led to the creation of policy, standards, and regulations allowing BlackBerry devices into many Limited Access areas at the Laboratory; and ensuring compliance with NNSA's classified and unclassified computer and telecommunications security requirements.

Progress in 2010

The Cyber Security Program (CSP) created a Security Operations Center (SOC) to answer LLNL cyber users' inquiries and issues and support the response capabilities of the cyber incident response team. The SOC is a Tier-1 and Tier-2 support center that monitors malicious activity on LLNL networks and system components. In addition, it supports a new Semantic Endpoint Protection (SEP) antivirus software that was installed on Windows and Macintosh computers. The upgraded software combines antivirus and advanced threat protection and integrates antivirus and intrusion-detection-tracking capabilities. With this improved detection capability, cyber security staff can quickly identify contaminated computer systems before data is lost or damaged. The SEP software, combined with SOC's enhanced monitoring, provides the Laboratory a stronger network defense.

CSP worked with the CIO and other Laboratory programs to complete the risk assessment and define standards, rules, and regulations for allowing LLNL-owned and managed BlackBerry devices into many Limited Access areas around the site. This project represented a paradigm shift away from the old policies that strictly prohibited all mobile communication devices in Limited Access areas. The process included obtaining approval from NNSA's Livermore Site Office to revise the institution's Emission Security (TEMPEST) protection plan and threat assessment, the two documents that address the proliferation of wireless devices and required precautions in Limited Access buildings. Staff worked directly with various work groups to plan, communicate, and coordinate the implementation of the new rules. CSP staff designed new signs and created a strategy to consistently post them across the Laboratory. The project significantly enhanced mobile computing capabilities at LLNL and improved employee productivity.

Cyber security staff coordinated several additional policy updates in 2010, including one that mandates the full-disk encryption of all LLNL laptops and mobile devices containing Unclassified Controlled Information that may leave the site. This policy was written to comply with NNSA requirements. A working group consisting of representatives from CSP, IT, Procurement, and major LLNL programs was established to develop and publish the policy and implementation plan.

Other policies that were updated include: an interim policy for embedded radio frequency identification readers, a policy for the Laptops on Foreign Travel pool, a new EnCase computer forensics procedure; and a new policy for microphones and cameras.

New NNSA cyber security requirements were implemented for the Laboratory's unclassified information systems, in accordance with the plan developed in 2009. Cyber security staff spent countless hours testing LLNL's information systems, which in turn led to improvements in the patching practices and configuration management of desktop systems. These changes also led to the development of 40 core protection services, including the management of LLNL's Active Directory servers and the institutional intrusion and prevention systems. The new core services now fall under the purview of service managers and are funded by the CIO and CSP.

A new cyber security policy mandates the full-disk encryption of Laboratory laptops and portable devices to safeguard unclassified controlled information.

Appendices

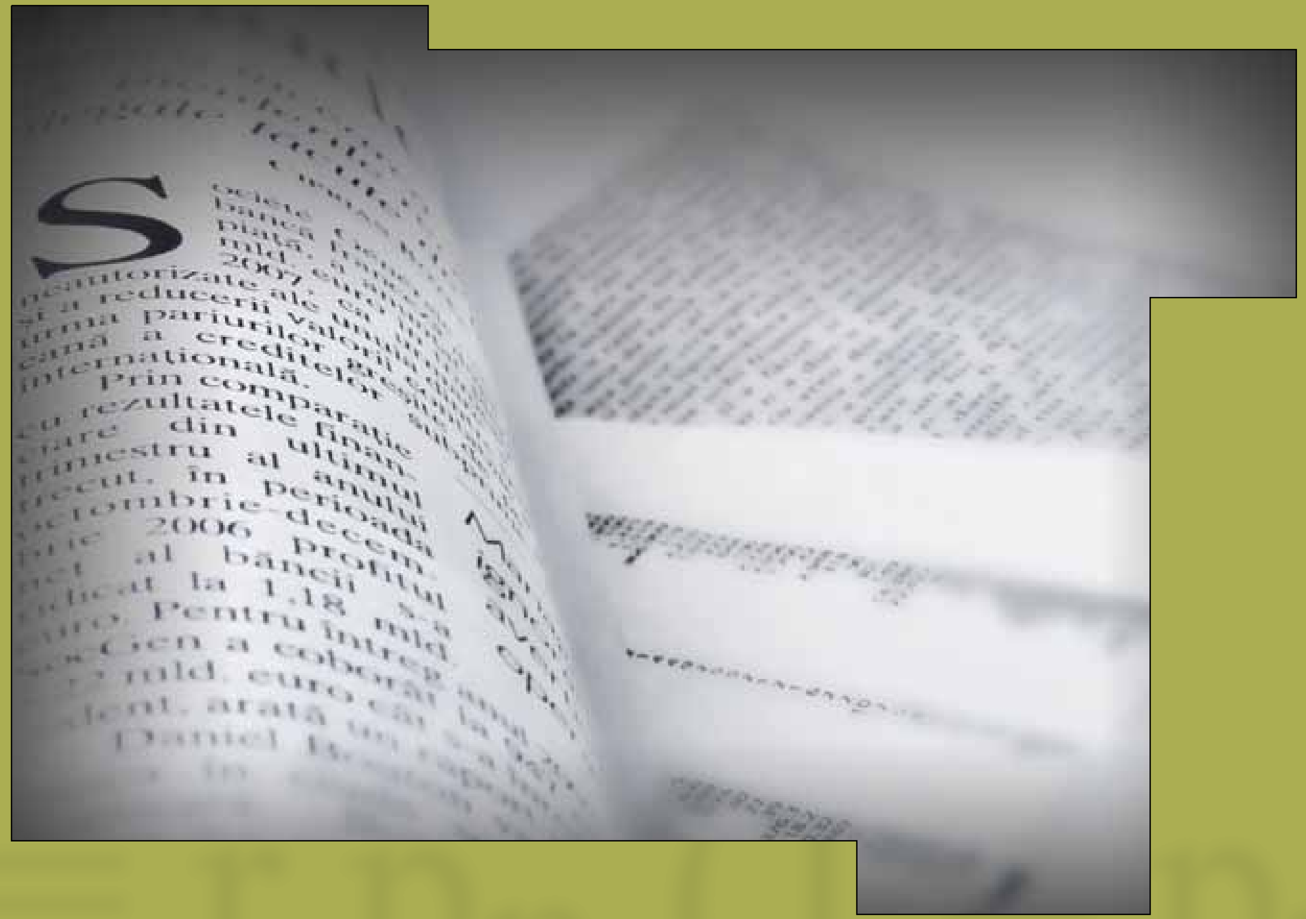

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Ball State University	Irene Livshits	Subcontract	Algebraic multigrid algorithms for finding many eigenpairs of partial differential operators	ASCR	Robert Falgout
Brigham Young University	Bryan Morse	Joint Research	Mosaics and super-resolution of unmanned aerial vehicle-acquired video using locally adaptive warping	LDRD	Mark Duchaineau and Jon Cohen
California Institute of Technology	Michael Ortiz	ASC Predictive Science Academic Alliance Program Center	Center for the Predictive Modeling and Simulation of High-Energy Density Dynamic Response of Materials	ASC	Dick Watson
California Polytechnic State University, San Luis Obispo	David Clague	Joint Research	Low-cost microarrays	LDRD	Tom Slezak
Cambridge University	Nikos Nikiforkis	Joint Research	Simulation and modeling using Overture	ASCR Base	Bill Henshaw
Carnegie Mellon University	Christos Faloutsos	Subcontract	Mining large dynamic weighted graphs	LDRD	Tina Eliassi-Rad
Carnegie Mellon University	Hanghang Tong	Collaboration	Mining large dynamic weighted graphs	LDRD	Tina Eliassi-Rad
Chalmers University	Sally McKee	Collaboration	Leveraging OpenAnalysis for alias analysis within ROSE	ASC	Dan Quinlan
Colorado State University	Michelle Strout and Sanjay Rajopadhye	Collaboration	Program analysis	ASCR	Dan Quinlan
Colorado State University	Donald Estep	Subcontract	A posteriori error analysis for hydrodynamic systems	LDRD	Carol Woodward
Columbia University	David Keyes	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
Columbia University	Ian Lipkin	Joint Research	Viral discovery and microarrays	DTRA	Tom Slezak
Cornell University	Claire Cardie	Joint Research	Coreference resolution	DHS	David Buttler

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Dresden University of Technology	Matthias Mueller	Joint Research	Semantic debugging of Message Passing Interface applications	ASCR, ASC, LDRD	Bronis de Supinski
Duke University	Herbert Edelsbrunner	Joint Research	Discrete methods for computing continuous functions	LDRD	Timo Bremer
Électricité de France	Phillippe Lafon	Collaboration	Aeroacoustics		Bill Henshaw
Emory University	Joel Saltz and Tahsin Kurc	Collaboration	Data science algorithms		Ghaleb Abdulla
FORTH-ICS and University of Crete	Demitrios Nikolopoulos	Joint Research	Power-aware computing for OpenMP programs	ASCR, ASC	Bronis de Supinski
Imperial College	Paul Kelly and Jose Gabriel de Figueiredo Coutinho	Collaboration	Field-programmable gate arrays research	ASCR	Dan Quinlan
Indiana University	Jeremiah Wilcock	Joint Research	Binary analysis	ASCR	Dan Quinlan
Johns Hopkins University	Allan Boyles	Collaboration	Seismoacoustic modeling for defense-related efforts		Shawn Larsen
Krell Institute	Lucille Kilmer	Subcontract	Department of Energy High-Performance Computer Science Fellowship Program: LLNL portion	ASC	John May
North Carolina State University	Frank Mueller	Joint Research	Compressing Message Passing Interface traces	ASCR, ASC	Bronis de Supinski
North Carolina State University	Vincent Freeh	Joint Research	Power-aware computing for Message Passing Interface programs	ASCR, ASC	Bronis de Supinski
Northern Arizona University	Paul Keim	Joint Research	Microbial forensics	DTRA	Tom Slezak
Ohio State University	P. Sadayappan and Christophe Alias	Collaboration	Optimizing compiler program analysis	ASCR	Dan Quinlan
Ohio University	Yusu Wang	Joint Research	Analysis and visualization of high-dimensional function	LDRD	Timo Bremer
Pennsylvania State University	Jinchao Xu, James Brannick, and Ludmil Zikatanov	Subcontract	Multigrid methods for numerical models arising in plasma simulations and quantum field theories	ASCR	Robert Falgout

APPENDIX A

COMPUTATION DIRECTORATE ANNUAL REPORT 2010

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Princeton University	Adam Burrows and Jason Nordhaus	Joint Research	Computational Astrophysics Consortium	ASCR SciDAC	Louis Howell
Purdue University	Ziqiang Cai	Summer Faculty	A posteriori error estimates for partial differential equations	ASC	Robert Falgout
Purdue University	Jayathi Murthy	ASC Predictive Science Academic Alliance Program Center	Center for Prediction of Reliability, Integrity, and Survivability of Microsystems (PRISM)	ASC	Dick Watson
Rensselaer Polytechnic Institute	Mark Shephard	Joint Research	Terascale simulation tools and technologies	ASCR SciDAC	Lori Diachin
Rensselaer Polytechnic Institute	Don Schwendeman	Subcontract	Development of numerical methods for mathematical models of high-speed reactive and nonreactive flow	ASCR Base	Bill Henshaw
Rice University	John Mellor-Crummey	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
Rice University	John Mellor-Crummey, Keith Cooper, and Vivek Sarkar	Collaboration	Use of ROSE for compiler optimizations	ASCR	Dan Quinlan
Royal Institute of Technology	Heinz-Otto Kreiss	Consultant	Adaptive methods for partial differential equations	ASCR Base	Lori Diachin and Anders Petersson
Southern Methodist University	Thomas Hagstrom	Joint Research	High-order structure grid methods for wave propagation on complex unbounded domains	ASCR Base	Bill Henshaw
Southern Methodist University	Dan Reynolds	Joint Research	Implicit solvers and preconditioning techniques for simulations of magnetohydrodynamics	ASCR SciDAC	Carol Woodward
Stanford University	Olav Lindtjorn	Collaboration	Reverse-time seismic imaging for hydrocarbon exploration	CRADA	Shawn Larsen

APPENDIX A

ACADEMIC OUTREACH

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Stanford University	Parvis Moin	ASC Predictive Science Academic Alliance Program Center	Center for Predictive Simulations of Multi-physics Flow Phenomena with Application to Integrated Hypersonic Systems	ASC	Dick Watson
Stanford University/ SLAC National Accelerator Center	Stuart Marshall and Jacek Becla	Joint Research	Large Synoptic Survey Telescope camera data output hardware and imaging formats	LDRD	Celeste Matarazzo
State University of New York, Stony Brook	Jim Glimm	Joint Research	Terascale simulation tools and technologies	ASCR SciDAC	Lori Diachin
State University of New York, Stony Brook	Xiao-Lin Li	Joint Research	Interoperable technologies for advanced petascale simulation	ASCR SciDAC	Lori Diachin
Swiss Federal Institute of Technology	Stephan Brunner	Subcontract	Efficient numerical algorithms for Vlasov simulation of laser–plasma interactions	LDRD	Jeff Hittinger
Technical University of Denmark	Robert Read	Collaboration	Water waves and wave energy generation		Bill Henshaw
Technical University of Vienna	Markus Schordan	Collaboration	Compiler construction	ASCR	Dan Quinlan
Texas A&M University	Marv Adams	Joint Research	Numerical methods for radiation transport	ASC	Peter Brown
Texas A&M University	Nancy Amato	Collaboration, Lawrence Scholar Program	Novel mechanisms to understand and improve load balance in Message Passing Interface applications	UCOP	Bronis de Supinski
Texas A&M University	Nancy Amato	Collaboration, Lawrence Scholar Program	Parallel graph algorithms	UCOP	Maya Gokhale
Texas A&M University	Bjarne Stroustrup and Lawrence Rauchwerger	Joint Research	Compiler construction and parallel optimizations	ASCR	Dan Quinlan

APPENDIX A

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Texas A&M University	Raytcho Lazarov and Yalchin Efendiev	Joint Research	Algebraic multigrid for Brinkman problems	ASCR Base	Panayot Vassilevski
Texas State University	Byron Gao	Collaboration	Search user interfaces; clustering	LDRD	David Buttler
Tufts University	Scott MacLachlan	Joint Research	Algebraic multigrid algorithms	ASC	Robert Falgout
UC Berkeley	James Demmel	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
UC Berkeley	Doug Dreger	Subcontract	Earthquake hazard	IGPP	Shawn Larsen
UC Berkeley	Kurt Miller and Michael Jordan	Collaboration, Lawrence Scholar Program	Latent variable models	LDRD	Tina Eliassi-Rad
UC Davis	Francois Gygi	Subcontract	Algorithms for electronic structure and first-principles molecular dynamics simulations using large-scale parallel computers	ASC	Art Mirin
UC Davis	Dave Wittman	Joint Research	Large Synoptic Survey Telescope long-term data analysis	LDRD	Celeste Matarazzo
UC Davis	Zhendong Su	Subcontract	Software security analysis research	LDRD	Dan Quinlan
UC Davis	Ken Joy	Joint Research	Visualization and Analytics Center for Enabling Technologies	ASCR SciDAC	Mark Duchaineau
UC Davis	Ken Joy	Subcontract	Improving accuracy and efficiency of 3D aerial video	DoD	Mark Duchaineau
UC Davis	Ken Joy	Collaboration, Lawrence Scholar Program	Discrete multi-material interface reconstruction for volume fraction data	UCOP	Mark Duchaineau
UC Davis	Ken Joy	Collaboration, Lawrence Scholar Program	Video processing research for the VidCharts and Persistics projects	UCOP	Mark Duchaineau

APPENDIX A

COMPUTATION DIRECTORATE ANNUAL REPORT 2010

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
UC Davis	Nelson Max	Joint Research	3D from wide-area aerial video	DOE Nonproliferation	Mark Duchaineau
UC Davis	Kwan-Liu Ma	Subcontract	Visualization of RADAR tomography reconstruction	DoD	Mark Duchaineau
UC Davis	Kwan-Liu Ma	Subcontract	Interactive tomographic reconstruction of aerial motion imagery	DoD	Mark Duchaineau
UC Davis	Chen-Nee Chuah	Collaboration	Network security	LDRD	Maya Gokhale
UC Davis	Kwan-Liu Ma	Joint Research	Large-graph data visualization	LDRD	Timo Bremer
UC Davis	Bernd Hamann	Joint Research	Analysis and visualization of scientific data using topology-based methods	LDRD	Timo Bremer
UC Davis	Michael Oskin	Joint Research	Topological analysis of geological data	LDRD	Timo Bremer
UC Los Angeles	Alfonso Cardenas	Collaboration, Lawrence Scholar Program	iScore: measuring the interestingness of articles in a limited user environment	UCOP	David Buttler
UC Los Angeles and Weizman Institute of Science	Achiezer Brandt	Consultant	Geometric and algebraic multigrid techniques	ASCR Base	Robert Falgout
UC Merced	Qinghua Guo	Joint Research	Learning from presence-only data	UCOP	Tina Eliassi-Rad
UC Riverside	Michalis Faloutsos	Collaboration	Behavioral approaches to traffic classification	LDRD	Tina Eliassi-Rad
UC Riverside	Marios Iliofotou	Collaboration	Behavioral approaches to traffic classification	LDRD	Tina Eliassi-Rad
UC San Diego	Allan Snavely and Laura Carrington	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
UC San Diego	Falko Kuester	Joint Research	Large-scale atomistic simulation visualization	ASC	Mark Duchaineau
UC San Diego	Falko Kuester	Joint Research	3D from video	DoD	Mark Duchaineau

APPENDIX A

COMPUTATION DIRECTORATE ANNUAL REPORT 2010

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
UC San Diego	Charles Elkan	Collaboration, Lawrence Scholar Program	Learning from presence-only data	UCOP	Tina Eliassi-Rad
UC San Diego	Randy Bank	Consultant	Numerical solutions of partial differential equations, multi-level iterative methods, and adaptive grid generations	ASC	Robert Falgout and Panayot Vassilevski
UC San Diego	Steve Swanson	Collaboration	Persistent memory emulator		Maya Gokhale
UC San Diego	Scott Baden	Subcontract	Data-driven execution of latency-tolerant algorithms	LLNL Overhead	Dan Quinlan
UC San Diego, Scripps Institution of Oceanography	Julie McClean	Collaboration	Ultra-high-resolution coupled climate simulations	BER	Art Mirin
UC Santa Barbara	Fred Chong and Diana Franklin	Joint Research	Reduction of memory footprints	ASC, LDRD	Bronis de Supinski
UC Santa Cruz	Carlos Maltzahn	Collaboration, Lawrence Scholar Program	Semantic file systems	LDRD	Maya Gokhale
UC Santa Cruz	Stan Woosley	Joint Research	Computational Astrophysics Consortium	ASCR SciDAC	Louis Howell
University of Arizona	David Lowenthal	Joint Research	Power-aware computing for Message Passing Interface programs; scalable performance models	ASCR, ASC	Bronis de Supinski
University of Arizona	Tim Axelrod	Joint Research	Large Synoptic Survey Telescope data management and high-speed data acquisition and analysis	LDRD	Celeste Matarazzo
University of British Columbia	Carl Ollivier-Gooch	Subcontract	Interoperable technologies for advanced petascale simulation	ASCR SciDAC	Lori Diachin

APPENDIX A

ACADEMIC OUTREACH

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
University of Chicago	Don Lamb	ASC Academic Strategic Alliance Program Center	Center for Astrophysical Thermonuclear Flashes	ASC	Dick Watson
University of Colorado	Ken Jansen	Joint Research	Interoperable technologies for advanced petascale simulation	ASCR SciDAC	Lori Diachin
University of Colorado	Tom Manteuffel	Joint Research	Solution methods for transport problems	ASC	Peter Brown
University of Colorado	Tom Manteuffel	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
University of Colorado	Steve McCormick	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
University of Colorado	Steve McCormick, Tom Manteuffel, John Ruge, Marian Brezina, Minho Park, Geoff Sanders, and Christian Ketelsen	Subcontract	Geometric and algebraic multigrid methods for quantum chromodynamics, magnetohydrodynamics, elasticity, transport, and other DOE applications	ASCR, ASC	Robert Falgout
University of Delaware	Richard Braun	Collaboration	Models of the eye	ASCR Base	Bill Henshaw
University of Houston	Yuriy Fofanov	Joint Research	Genomic algorithms	DHS	Tom Slezak
University of Illinois, Urbana-Champaign	Michael Heath	ASC Academic Strategic Alliance Program Center	Center for Simulation of Advanced Rockets	ASC	Dick Watson
University of Illinois/ IBM	Hormozd Gahvari and Kirk Jordon	Collaboration	Modeling algebraic multigrid performance on multicore architectures	LDRD	Ulrike Yang
University of Illinois/ National Center of Supercomputing Applications	Ray Plante	Joint Research	Large Synoptic Survey Telescope data management and high-speed data acquisition and analysis	LDRD	Celeste Matarazzo
University of Louisville	Yongsheng Liam	Collaboration	Micro-air vehicles		Bill Henshaw
University of Maryland	Jeff Hollingsworth	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski

APPENDIX A

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
University of Massachusetts, Amherst	Andrew McCallum	Joint Research	Cross-language topic models	LDRD	David Buttler
University of Michigan	R. Paul Drake	ASC Predictive Science Academic Alliance Program Center	Center for Radiative Shock Hydrodynamics (CRASH)	ASC	Dick Watson
University of Munich	Dieter August Kranzlmueller	Joint Research	Detecting communication patterns to optimize applications	ASCR, ASC	Bronis de Supinski
University of Nevada, Reno	John Louie	Collaboration	Seismic modeling in the basin and range region		Shawn Larsen
University of New Mexico	Dorian Arnold	Joint Research	Scalable tool infrastructures	ASC, ASCR	Bronis de Supinski
University of North Carolina	Robert Fowler	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of North Carolina	Jan Prins	Joint Research	Efficient OpenMP runtimes for tasking	ASC, LDRD	Bronis de Supinski
University of Southern California	Robert Lucas, Mary Hall, and Jacqueline Chame	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of Southern California	Gerard Medioni	Subcontract	Activity analysis in wide-area overhead video	DOE Nonproliferation	Mark Duchaineau
University of Tennessee	Jack Dongarra and Shirley Moore	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of Tennessee	Jack Dongarra	Joint Research	Empirical tuning	ASCR	Dan Quinlan
University of Texas, Austin	Omar Ghattas	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
University of Texas, Austin	Robert Moser	ASC Predictive Science Academic Alliance Program Center	Center for Predictive Engineering and Computational Sciences (PECOS)	ASC	Dick Watson

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
University of Texas, San Antonio	Qing Yi	Subcontract	Program analysis and optimization for the empirical tuning of scientific applications	ASCR	Dan Quinlan
University of Utah	Valerio Pascucci	Subcontract	Algorithms for geometric processing, image segmentation, and data streaming	LDRD	Timo Bremer
University of Utah	Ross Whitaker	Joint Research	Analysis and visualization of high-dimensional function	LDRD	Timo Bremer
University of Utah	Ellen Riloff	Subcontract	Coreference resolution research	LDRD	David Buttler
University of Utah	Ganesh Gopalakrishnan	Joint Research	Semantic debugging of Message Passing Interface applications	ASCR, ASC, LDRD	Bronis de Supinski
University of Utah	Charles Hansen	Consultant	Data exploration, multiresolution scientific data visualization, and algorithm design	ASC CSSE	Mark Duchaineau
University of Utah	Chris Johnson, Valerio Pascucci, Chuck Hansen, Claudio Silva, Lee Myers, Allen Sanderson, and Steve Parker	Joint Research	Visualization and Analytics Center for Enabling Technologies	ASCR SciDAC	Mark Duchaineau
University of Utah	Gannesh Gopalakrishnan	Collaboration	Message Passing Interface optimizations	ASCR	Dan Quinlan
University of Utah	David Pershing	ASC Academic Strategic Alliance Program Center	Center for the Simulation of Accidental Fires and Explosions	ASC	Dick Watson
University of Washington	Carl Ebeling and Scott Hauck	Collaboration	Coarse-grain processor architectures		Maya Gokhale
University of Waterloo	Hans de Sterck	Joint Research	Algebraic multigrid for Markov chains	ASCR SciDAC	Ulrike Yang

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
University of Wisconsin	Bart Miller	Joint Research	Scalable debugging	ASCR, ASC, LRD	Bronis de Supinski
University of Wisconsin	Ben Liblit	Joint Research	Scalable debugging	ASCR, ASC	Bronis de Supinski
University of Wisconsin	Karu Sankaralingam	Joint Research	Resilient computing	ASCR, ASC, LDRD	Bronis de Supinski
University of Wisconsin	Jason Kraftcheck	Joint Research	Mesquite software development	ASCR Base	Lori Diachin
University of Wisconsin	Bart Miller	Joint Research	Support for enhanced dyninst testing and initial steps toward open binary editing environment	ASC CSSE	Martin Schulz
University of Wuppertal	Karsten Kahl	Visiting Researcher	Algebraic multigrid algorithms	ASC	Robert Falgout
Virginia Institute of Technology	Kirk Cameron	Joint Research	Power-aware computing for hybrid systems	ASCR, ASC	Bronis de Supinski
Worcester Polytechnic University	Homer Walker	Joint Research	Nonlinear solvers and subsurface simulation methods	ASCR SciDAC	Carol Woodward

RETURN TO TOC

Book Chapters

de Supinski, B.R. and M. Schulz, "Flexible Tools Supporting a Scalable First-Principles MD Code," in *Scientific Computer Performance*, D.H. Bailey, R. F. Lucas, and S. Williams (Taylor and Francis, New York) (UCRL-JRNL-445511).

Schulz, M., "Checkpointing," in *Encyclopedia of Parallel Computing*, D. Padua (Springer, New York) (LLNL-MI-419844).

Slezak, T., S.N. Gardner, and J. Allen, "The Design of Genomic Signatures for Pathogen Identification and Characterization," in *Microbial Forensics*, 2nd edition, B. Budowle, et al. (Elsevier), 493–508.

Journal Papers

Aagaard, B.T., et al., "Ground Motion Modeling of Hayward Fault Scenario Earthquakes II: Simulation of Long-Period and Broadband Ground Motions," *B. Seismol. Soc. Am.* **100** (6), 2945–2977 (LLNL-JRNL-419883).

Adler, J.H., et al., "Nested Iteration and First-Order System Least Squares for Incompressible Resistive Magneto-Hydrodynamics," *SIAM J. Sci. Comput.* **32** (3), 1506–1526.

Allen, J., et al., "Performance of Prototypes for the ALICE Electromagnetic Calorimeter," *Nucl. Instrum. Meth. A* **615** (1), 6–13.

Allgaier, M., et al. "Targeted Discovery of Glycoside Hydrolases from a Switchgrass-Adapted Compost Community," *PLoS ONE* **5** (1).

Almgren, A.S., et al. "Castro: A New Compressible Astrophysical Solver. I. Hydrodynamics and Self-Gravity," *Astrophys. J.* **715** (2), 1221–1238.

Amendt, P., J. Milovich, L.J. Perkins, and H. Robey, "An Indirect-Drive Noncryogenic Double-Shell Path to 1ω Nd-Laser Hybrid Inertial Fusion–Fission Energy," *Nucl. Fusion* **50** (10).

Anderson, J.C., C. Garth, M.A. Duchaineau, and K.I. Joy, "Smooth, Volume-Accurate Material Interface Reconstruction," *IEEE T. Vis. Comput. Gr.* **16** (5), 802–814.

Appelo, D. and N.A. Petersson, "A Fourth Order Accurate Embedded Boundary Method for the Wave Equation," *SIAM J. Sci. Comput.* (LLNL-JRNL-417163).

Baker, A.H., T.V. Kolev, and U.M. Yang, "Improving Algebraic Multigrid Interpolation Operators for Linear Elasticity Problems," *Numer. Linear Algebr.* **17** (2–3), 495–517 (LLNL-JRNL-412928).

Banks, J.W., "On Exact Conservation for the Euler Equations with Complex Equations of State," *Commun. Comput. Phys.* **8** (5), 995–1015 (LLNL-JRNL-410514).

Banks, J.W. and J.A.F. Hittinger, "A New Class of Non-Linear, Finite-Volume Methods for Vlasov Simulation," *IEEE T. Plasma Sci.* **38** (9), 2198–2207 (LLNL-JRNL-420843).

Banks, J.W. and B. Sjogreen, "A Normal Mode Stability Analysis of Numerical Interface Conditions for Fluid/Structure Interaction," *Commun. Comput. Phys.* (LLNL-JRNL-415710).

Brain, D.A., et al., "Espisodic Detachment of Martian Crustal Magnetic Fields Leading to Bulk Atmospheric Plasma Escape," *Geophys. Res. Lett.* **37** (LLNL-JRNL-413528).

Brannick, J.J. and R.D. Falgout, "Compatible Relaxation and Coarsening in Algebraic Multigrid," *SIAM J. Sci. Comput.* **32** (2), 1393–1416.

Bremer, P.-T., et al., "Analyzing and Tracking Burning Structures in Lean Premixed Hydrogen Flames," *IEEE T. Vis. Comput. Gr.* **16** (2), 248–260.

Brezina, M., et al., "Towards Adaptive Smoothed Aggregation Multigrid (aSA) for Nonsymmetric Problems," *SIAM J. Sci. Comput.* **32** (1), 14–39.

Cai, Z.Q., C. Tong, P.S. Vassilevski, and C.B. Wang, "Mixed Finite Element Methods for Incompressible Flow: Stationary Stokes Equations," *Numer. Meth. Part. D. E.* **26** (4), 957–978.

Childs, H., et al., "Extreme Scaling of Production Visualization Software on Diverse Architectures," *IEEE Comput. Graph.* **30** (3), 22–31.

Colella, P., M.R. Dorr, J.A.F. Hittinger, and D.F. Martin, "High-Order, Finite-Volume Methods in Mapped Coordinates," *J. Comput. Phys.* (LLNL-JRNL-422807).

Correa, C.D. and K.-L. Ma, "Dynamic Video Narratives," *ACM T. Graphic.* **29** (4).

Correa, C.D., D. Silver, and M. Chen, "Constrained Illustrative Volume Deformation," *Computers and Graphics* **34** (4), 370–377.

De Sterck, H., et al., "Algebraic Multigrid for Markov Chains," *SIAM J. Sci. Comput.* **32** (2), 544–562.

De Sterck, H., et al., "Smoothed Aggregation Multigrid for Markov Chains," *SIAM J. Sci. Comput.* **32** (1), 40–61.

De Sterck, H., V.E. Henson, and G. Sanders, "Multilevel Aggregation Methods for Small-World Graphs with Application to Random-Walk Ranking," *Comput. Inform.*

De Sterck, H., K. Miller, G. Sanders, and M. Winlaw, "Recursively Accelerated Multilevel Aggregation for Markov Chains," *SIAM J. Sci. Comput.* **32** (3), 1652–1671.

De Sterck, H., K. Miller, T. Manteuffel, and G. Sanders, "Top-Level Acceleration of Adaptive Algebraic Multilevel Methods for Steady-State Solution to Markov Chains," *Adv. Comput. Math.* (1–29), 1019–7168.

DeAngelis, K.M., et al., "Strategies for Enhancing the Effectiveness of Metagenomic-Based Enzyme Discovery in Lignocellulolytic Microbial Communities," *Bioenergy Res.* **3** (2), 146–158.

Dobrev, V., J.L. Guermond, and B. Popov, "Surface Reconstruction and Image Enhancement via L-1-Minimization," *SIAM J. Sci. Comput.* **32** (2), 1591–1616.

Dorr, M.R., et al., "A Numerical Algorithm for the Solution of a Phase-Field Model of Polycrystalline Materials," *J. Comput. Phys.* **229** (3), 626–641 (LLNL-JRNL-409478).

Fattebert, J.L., "Accelerated Block Preconditioned Gradient Method for Large-Scale Wave Functions Calculations in Density Functional Theory," *J. Comput. Phys.* **229** (2), 441–452 (LLNL-JRNL-405947).

Gallagher, B. and T. Eliassi-Rad, "Leveraging Label-Independent Features for Classification in Sparsely Labeled Networks: An Empirical Study," *Lect. Notes Comput. Sc.* **5498**, 1–19 (LLNL-JRNL-411529).

Gardner, S.N., C.J. Jaing, K.S. McLoughlin, and T.R. Slezak, "A Microbial Detection Array (MDA) for Viral and Bacterial Detection," *BMC Genomics* **11**.

Gerber, S., P.-T. Bremer, V. Pascucci, and R. Whitaker, "Visual Exploration of High Dimensional Scalar Functions," *IEEE T. Vis. Comput. Gr.* **16** (6), 1271–1280 (LLNL-CONF-446831).

Hauser, J., et al., "Probabilistic Seismic 3D Model for the European Arctic," *J. Geophys. Res.*

Hazen, T.C., et al., "Deep-Sea Oil Plume Enriches Indigenous Oil-Degrading Bacteria," *Science* **330** (6001), 204–208.

Hoefler, T., et al., "The Scalable Process Topology Interface of MPI 2.2," *Concurr. Comp.-Pract. E.* (LLNL-JRNL-438133).

Inglis, T., et al., "Multilevel Space–Time Aggregation for Bright Field Cell Microscopy Segmentation and Tracking," *Int. J. Biomedical Imaging.*

Isenburg, M., P. Lindstrom, and H. Childs, "Parallel and Streaming Generation of Ghost Data for Structured Grids," *IEEE Comput. Graph.* **30** (3), 32–44 (UCRL-JRNL-403175).

Johnsen, E., et al., "Assessment of High-Resolution Methods for Numerical Simulations of Compressible Turbulence with Shock Waves," *J. Comput. Phys.* **229** (4), 1213–1237.

Khachatryan, V., et al., "Observation of Long-Range, Near-Side Angular Correlations in Proton–Proton Collisions at the LHC," *J. High Energy Phys.* **9**.

Kilbourne, K.H., et al., "Coral Windows Onto Seasonal Climate Variability in the Northern Caribbean Since 1479," *Geochem. Geophy. Geosy.* **11**.

Kim, A., D.S. Dreger, and S. Larsen, "Moderate Earthquake Ground-Motion Validation in the San Francisco Bay Area," *B. Seismol. Soc. Am.* **100** (2), 819–825.

Kollet, S.J., et al., "Proof of Concept of Regional Scale Hydrologic Simulations at Hydrologic Resolution Utilizing Massively Parallel Computer Resources," *Water Resour. Res.* **46** (LLNL-JRNL-462006).

Kreiss, H.O., O.E. Ortiz, and N.A. Petersson, "Initial-Boundary Value Problems for Second Order Systems of Partial Differential Equations," *Mathematical Modeling and Numerical Analysis* (LLNL-JRNL-416303).

Landen, O.L., et al., "Capsule Performance Optimization in the National Ignition Campaign," *Phys. Plasmas* **17** (5).

Liao, C.H., D.J. Quinlan, J.J. Willcock, and T. Panas, "Semantic-Aware Automatic Parallelization of Modern Applications Using High-Level Abstractions," *Int. J. Parallel Prog.* **38** (5–6), 361–378.

Maki, K.L., et al., "Tear Film Dynamics On an Eye-Shaped Domain II: Flux Boundary Conditions," *J. Fluid Mech.* **647**, 361–390.

Maki, K.L., R.J. Braun, W.D. Henshaw, and P.E. King-Smith, "Tear Film Dynamics On an Eye-Shaped Domain I: Pressure Boundary Conditions," *Math. Med. Biol.* **27** (3), 227–254.

Manteuffel, T.A., et al., "Further Results on Error Estimators for Local Refinement with First-Order System Least Squares (FOSLS)," *Numer. Linear Algebr.* **17** (2–3), 387–413.

Meezan, N.B., et al., "National Ignition Campaign Hohlraum Energetics," *Phys. Plasmas* **17** (5).

Neville, J., B. Gallagher, T. Eliassi-Rad, and T. Wang, "Correcting Evaluation Bias of Relational Classifiers with Network Cross Validation," *Knowl. Inf. Syst.* (LLNL-JRNL-455699).

Ollivier-Gooch, C., et al., "An Interoperable, Data-Structure-Neutral Component for Mesh Query and Manipulation," *ACM T. Math. Software* **37** (3).

Petersson, N.A. and B. Sjogreen, "Stable Grid Refinement and Singular Source Discretization for Seismic Wave Aimulations," *Commun. Comput. Phys.* **8** (5) 1074–1110 (LLNL-JRNL-419382).

Ping, Y., et al., "Warm Dense Matter Created by Isochoric Laser Heating," *High Energ. Dens. Phys.* **6** (2), 246–257.

Reynolds, D.R., R. Samtaney, and C.S. Woodward, "Operator-Based Preconditioning of Stiff Hyperbolic Systems," *SIAM J. Sci. Comput.* **32** (1), 150–170 (LLNL-JRNL-413456).

Robey, H.F., et al., "High Performance Capsule Implosions on the OMEGA Laser Facility with Rugby Hohlraums," *Phys. Plasmas* **17** (5).

Rodgers, A.J., N.A. Petersson, and B. Sjogreen, "Simulation of Topographic Effects on Seismic Waves from Shallow Explosions Near the North Korean Nuclear Test Site with Emphasis on Shear Wave Generation," *J. Geophys. Res.-Sol. Ea.* **115**.

Rübel, O., et al., "Integrating Data Clustering and Visualization for the Analysis of 3D Gene Expression Data," *IEEE-ACM T. Comput. Bi.* **7** (1), 64–79.

Schwendeman, D.W., A.K. Kapila, and W.D. Henshaw, "A Study of Detonation Diffraction and Failure for a Model of Compressible Two-Phase Reactive Flow," *Combust. Theor. Model.* **14** (3), 331–366.

Shadid, J.N., et al., "Towards a Scalable Fully Implicit Fully Coupled Resistive MHD Formulation with Stabilized FE Methods," *J. Comput. Phys.* **229** (20), 7649–7671.

Shelton, R.T., et al., "Target Diagnostic Control System Implementation for the National Ignition Facility," *Rev. Sci. Instrum.* **81** (10) (LLNL-CONF-43217).

Singer, M.A., S.L. Wang, and D.P. Diachin, "Design Optimization of Vena Cava Filters: An Application to Dual Filtration Devices," *J. Biomech. Eng.-T. ASME* **132** (10) (LLNL-JRNL-421025).

Stengle, G., J. McEnerney, and R. Robson, "Convex Polarities Over Ordered Fields," *J. Pure Appl. Algebra* **214** (4), 370–379.

Summa, B., et al., "Interactive Editing of Massive Imagery Made Simple: Turning Atlanta into Atlantis," *ACM T. Graphic.* **29** (5).

Tikhonova, A., C.D. Correa, and K.-L. Ma, "An Exploratory Technique for Coherent Visualization of Time-Varying Volume Data," *Comput. Graph. Forum* **29** (3), 783–792.

Tikhonova, A., C.D. Correa, and K.-L. Ma, "Visualization by Proxy: A Novel Framework for Deferred Interaction with Volume Data," *IEEE T. Vis. Comput. Gr.* **16** (6), 1551–1559.

Tong, C., "Self-Validated Variance-Based Methods for Sensitivity Analysis of Model Outputs," *Reliab. Eng. Syst. Safe.* **95** (3), 301–309.

Torres, C.L., et al., "LAVA: An Open-Source Approach To Designing LAMP (Loop-Mediated Isothermal Amplification) DNA Signatures," *BMC Bioinformatics*.

Ulmer, C., et al., "Massively Parallel Acceleration of a Document-Similarity Classifier to Detect Web Attacks," *J. Parallel Distr. Com.* **71** (2), 225–235 (LLNL-JRNL-422668).

Vassilevski, P.S., "General Constrained Energy Minimization Interpolation Mappings for AMG," *SIAM J. Sci. Comput.* **32** (1), 1–13.

Victoria, J.G., et al., "Viral Nucleic Acids in Live-Attenuated Vaccines: Detection of Minority Variants and an Adventitious Virus," *J. Virol.* **84** (12), 6033–6040 (LLNL-JRNL-426582).

Wang, W., et al., "Construction of Low Dissipative High Order Well-Balanced Filter Schemes for Non-Equilibrium Flows," *J. Comput. Phys.* (LLNL-JRNL-420355).

Wang, P., "Modeling Material Responses by Arbitrary Lagrangian Eulerian Formulation and Adaptive Mesh Refinement Method," *J. Comput. Phys.* **229** (5), 1573–1599.

Wang, S.L. and M.A. Singer, "Toward an Optimal Position for Inferior Vena Cava Filters: Computational Modeling of the Impact of Renal Vein Inflow with Celect and TrapEase Filters," *J. Vasc. Interven. Radiol.* **21** (3), 367–374 (LLNL-JRNL-414674).

Wehner, M.F., et al., "Towards Direct Simulation of Future Tropical Cyclong Statistics in a High Resolution Global Atmospheric Model," *Advances in Meteorology* **2010** (LLNL-JRNL-428972).

Wheeler, K.E., et al., "Functional Insights Through Structural Modeling of Proteins Expressed in a Natural Microbial Community," *J. Proteomics & Bioinformatics* **3** (9), 266–274 (LLNL-JRNL-421048).

Xie, S.C., et al., "Observed Large-Scale Structures and Diabatic Heating and Drying Profiles During TWP-ICE," *J. Climate* **23** (1), 57–79.

Xie, S.C., et al., "Arm Climate Modeling Best Estimate Data: A New Data Product for Climate Studies," *B. Am. Meteorol. Soc.* **91** (1).

Yang, U.M., "On Long-Range Interpolation Operators for Aggressive Coarsening," *Numer. Linear Algebr.* **17** (2–3), 453–472 (LLNL-JRNL-417371).

Chen, Y., et al., "A Method to Detect Differential Gene Expression," *Biomedical Informatics Insights* (LLNL-JRNL-421545).

Conference Proceedings

Ahn, D.H., et al., "Scalable Temporal Order Analysis for Large-Scale Debugging," *SC2009*, (ACM Int. Conf. Supercomputing), (LLNL-PROC-412227).

Bailey, T.S. and R.D. Falgout, "Analysis of Massively Parallel Discrete-Ordinates Transport Sweep Algorithms with Collisions," *International Conference on Mathematics, Computational Methods, and Reactor Physics*, (M&C 2009), (LLNL-CONF-407968).

Baker, A.H., M. Schulz, and U.M. Yang, "On the Performance of an Algebraic Multigrid Solver on Multicore Clusters," *VECPAR 2010*, (LLNL-CONF-429864).

Baldwin, C., G. Abdulla, and T. Critchlow, "Multi-Resolution Modeling of Large-Scale Scientific Simulation Data," *12th International Conference on Information and Knowledge Management*.

Barnes, B., et al., "Using Focused Regression for Accurate Time-Constrained Scaling of Scientific Applications," *24th International Parallel and Distributed Processing Symposium (IPDPS 2010)*, (IEEE), (LLNL-CONF-422989).

Bihari, B.L., "Applicability of Transactional Memory to Modern Codes," *International Conference on Numerical Analysis and Applied Mathematics*, (American Institute of Physics), **1281**, (LLNL-PROC-444631).

Biswas, S., et al., "PSMalloc: Content-Based Memory Management for MPI Applications," *MEDEA 2009* (LLNL-PROC-414508).

Bronevetsky, G., et al., "AutomaDeD: Automata-Based Debugging for Dissimilar Parallel Tasks," *40th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2010)*, (IEEE and IFIP), (LLNL-CONF-426270).

Bronevetsky, G., et al., "Statistical Fault Detection for Parallel Applications with AutomaDeD," *IEEE Workshop on Silicon Errors in Logic: System Effects*.

de Callafon, R., J. Meyer, and N. Delson, "Design, Modeling, and Stabilization of a Moment Exchange Based Inverted Pendulum," *15th IFAC Symposium on System Identification*, **15**.

De Sterck, H., et al., "Fast Multilevel Methods for Markov Chains," *6th International Workshop on the Numerical Solution of Markov Chains*, (Numerical Solution of Markov Chain).

De Sterck, H., V.E. Henson, and G.D. Sanders, "Multilevel Aggregation Methods for Small-World Graphs with Application to Random-Walk Ranking," *Workshop on Information Retrieval by Matrix Methods in the International Supercomputing Conference*, (ACM).

Dorr, M.R., et al., "Numerical Simulation of Phase Space Advection in Gyrokinetic Models of Fusion Plasmas," *SciDAC 2010*, (LLNL-PROC-438645).

Eliassi-Rad, T. and K. Henderson, "Literature Search Through Mixed-Membership Community Discovery," *3rd International Workshop on Social Computing, Behavioral Modeling, and Prediction*, (Lect. Notes Comput. Sc., Springer–Verlag Berlin), **6007**, 70–78.

Fattebert, J.-L., "Nonlinear Eigenvalue Problems in Density Functional Theory Calculations," *Mathematisches Forschungsinstitut Oberwolfach: Linear and Nonlinear Eigenproblems for PDEs*, (LLNL-PROC-416304).

Frazier, T., et al., "The National Ignition Facility Data Repository," *2009 International Conference on Accelerators and Large Experimental Physics Control Systems*, (LLNL-CONF-417392).

Gallagher, B. and T. Eliassi-Rad, "Leveraging Label-Independent Features for Classification in Sparsely Labeled Networks: An Empirical Study," *2nd International Workshop on Social Network Analysis (SNAKDD 2008)*, (Lect. Notes Comput. Sc., Springer–Verlag Berlin), **5498**, 1–19.

Gallagher, B., M. Iliofotou, T. Eliassi-Rad, and M. Faloutsos, "Homophily in Application Layer and Its Usage in Traffic Classification," *IEEE Conference on Computer Communications (INFOCOM)*.

PUBLICATIONS

Gamblin, T., et al., "Clustering Performance Data Efficiently at Massive Scales," *24th International Conference on Supercomputing (ICS 2010)*, (ACM), (LLNL-CONF-422684).

Goodell, D., et al., "Minimizing MPI Resource Contention in Multithreaded Multicore Environments," *IEEE Conference on Cluster Computing 2010 (Cluster 2010)*, (IEEE), (LLNL-CONF-442477).

Gu, Y., L. Lu, R. Grossman, and A. Yoo, "Processing Massive-Sized Graphs Using Sector/Sphere," *3rd IEEE Workshop on Many-Task Computing on Grids and Supercomputers*.

Hilbrich, T., M. Schulz, B.R. de Supinski, and M.S. Mueller, "MUST: A Scalable Approach to Runtime Error Detection in MPI Programs," *3rd Parallel Tools Workshop*.

Hoefler, T., et al., "Efficient MPI Support for Advanced Hybrid Programming Models," *EuroMPI 2010*, (LLNL-CONF-434306).

Humphrey, A., et al., "Verification for Portability, Scalability, and Grokkability," *Workshop on Exploiting Concurrency Efficiently and Correctly*.

Iliofotou, M., et al., "Profiling-by-Association: A Resilient Traffic Profiling Solution for the Internet Backbone," *ACM International Conference on Emerging Networking EXperiments and Technologies (CoNEXT)*, (LLNL-CONF-461043).

Izumi, N., et al., "Experimental Study of Neutron Induced Background Noise on Gated X-Ray Framing Cameras," *Topical Conference on High-Temperature Plasma Diagnostics*, (Rev. Sci. Instrum., AIP), **81**, 10E515.

Jain, S. et al., "Extracting the Textual and Temporal Structure of Supercomputing Logs," *International Conference on High Performance Computing*, (LLNL-CONF-416291).

Kaplan, I., "Evaluating Use of Data Flow Systems for Large Graph Analysis," *2nd Workshop on Many-Task Computing on Grids and Supercomputers*.

Kettimuthu, R., et al., "Lessons Learned from Moving Earth System Grid Data Sets Over a 20 Gbps Wide-Area Network," *ACM International Symposium on High Performance Distributed Computing (HPDC)*, (International Conference for High Performance Computing, Networking, Storage, and Analysis, OmniPress).

Koch, J., et al., "A Hardened Gated X-Ray Imaging Diagnostic for Inertial Confinement Fusion Experiments at the National Ignition Facility," *Topical Conference on High-Temperature Plasma Diagnostics*, (Rev. Sci. Instrum., AIP), **81**, 10E539.

Koniges, A., et al., "Application Acceleration on Current and Future Cray Platforms," *Cray Users Group 2010: Simulation Comes of Age*, (Cray Users Group Proceedings).

Kyrala, G.A., et al., "Measuring Symmetry of Implosions in Cryogenic Hohlraums at NIF Using Gated X-Ray Detectors," *Topical Conference on High-Temperature Plasma Diagnostics*, (Rev. Sci. Instrum., AIP), **81**, 10E316.

Landen, O.L., et al., "Capsule Performance Optimization in the National Ignition Campaign," *51st Annual Meeting of the Division of Plasma Physics of the American Physics Society*, (Phys. Plasmas, AIP), **17**, 056301.

Li, D., et al., "Hybrid MPI/OpenMP Power-Aware Computing," *24th International Parallel and Distributed Processing Symposium (IPDPS 2010)*, (IEEE), (LLNL-CONF-422990).

Li, D., et al., "Power-Aware MPI Task Aggregation Prediction for High-End Computing Systems," *24th International Parallel and Distributed Processing Symposium (IPDPS 2010)*, (IEEE), (LLNL-CONF-422991).

Liao, C.H., D.J. Quinlan, T. Panas, and B.R. de Supinski, "A ROSE-Based OpenMP 3.0 Research Compiler Supporting Multiple Runtime Libraries," *6th International Workshop on OpenMP 2010*, (Lect. Notes Comput. Sc., Springer–Verlag Berlin), (LLNL-CONF-422873).

Liao, C.H., D.J. Quinlan, R. Vuduc, and T. Panas, "Effective Source-to-Source Outlining to Support Whole Program Empirical Optimization," *22nd International Workshop on Languages and Compilers for Parallel Computing*, (Lect. Notes Comput. Sc., Springer–Verlag Berlin), **5898**, 308–322, (LLNL-CONF-414916).

Lindstrom P., "Triangulation Made Easy," *IEEE Conference on Computer Vision and Pattern Recognition 2010*, (Computer Vision and Pattern Recognition, IEEE), 1554–1561, (LLNL-CONF-422879).

Lindstrom, P. and J. Cohen, "On-the-Fly Decompression and Rendering of Multiresolution Terrain," *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games 2010*, (ACM), (LLNL-CONF-411859).

Meezan, N.B., et al., "National Ignition Campaign Hohlraum Energetics," *51st Annual Meeting of the Division of Plasma Physics of the American Physics Society*, (Phys. Plasmas, AIP), **17**.

Moody, A., G. Bronevetsky, K. Mohror, and B.R. de Supinski, "Modeling and Design of a Scalable Multilevel Checkpoint System," *SC10*, (IEEE and ACM), (LLNL-CONF-427742).

Mueller, F., et al., "ScalaTrace: Tracing, Analysis, and Modeling of HPC Codes at Scale," *Para 2010: State of the Art in Scientific and Parallel Computing*, (Springer), (LLNL-CONF-427005).

Parker, W.E. and Y.M. Dardenne, "Reanalysis of Collected Test Debris (U)," *ILWOG 42 Interlaboratory Working Group for Stockpile Radiochemistry*, (Proceedings of ILWOG 42 2010, Los Alamos National Laboratory), **42**.

Parker, W.E., et al., "Radiochemistry Detectors and Primary Performance (U)," *ILWOG 42 Interlaboratory Working Group for Stockpile Radiochemistry*, (Proceedings of ILWOG 42 2010, Lawrence Livermore National Laboratory).

Ping, Y., et al., "Warm Dense Matter Created by Isochoric Laser Heating," *2nd International Conference on High Energy Density Physics*, (High Energ. Dens. Phys., Elsevier Science BV), **6**, 246–257.

Pirkelbauer, P., D. Dechev, and B. Stroustrup, "Support for the Evolution of C++ Generic Functions," *Software Language Engineering*, (Lect. Notes Comput. Sc., Springer–Verlag).

Preissl, R., et al., "Exploitation of Dynamic Communication Patterns Through Static Analysis," *2010 International Conference on Parallel Processing (ICPP-10)*, (IACC), (LLNL-CONF-438991).

Reynolds, D.R., R. Samtaney, and C.S. Woodward, "Operator-Based Preconditioning of Stiff Hyperbolic Systems," *10th Copper Mountain Conference on Iterative Methods*, (SIAM J. Sci. Comput., SIAM Publications), **32**, 150-170.

Richards, D.F., et al., "Beyond Homogeneous Decomposition: Scaling Long-Range Forces on Massively Parallel Architectures," *SC09*, (ACM), (LLNL-CONF-411834).

Robey, H.F., et al., "High Performance Capsule Implosions on the OMEGA Laser Facility with Rugby Hohlraums," *51st Annual Meeting of the Division of Plasma Physics of the American Physics Society*, (Phys. Plasmas, AIP), 17.

Rübel, O., et al., "Coupling Visualization and Data Analysis for Knowledge Discovery from Multidimensional Scientific Data," *International Conference on Computational Science (ICCS 2010)*, (Procedia Computer Science, Elsevier), **1**.

Rübel, O., et al., "Linking Advanced Visualization and MATLAB for the Analysis of 3D Gene Expression Data," *2nd International Workshop on Visualization in Medicine and Life Sciences*, (Springer).

Schulz, M., et al., "On the Performance of the Miranda CFD Code on Multicore Architectures," *Parallel CFD*, (PERI), **78**, (LLNL-ABS-411404).

Schulz, M., et al., "Analyzing the Performance of Scientific Applications with Open|SpeedShop," *21st International Conference on Parallel Computational Fluid Dynamics (ParallelCFD 2009)*, (LLNL-CONF-418135).

Singh, K., "Comparing Scalability Prediction Strategies on an SMP of CMPs," *Euro-Par 2010*, (LLNL-CONF-423717).

Sjogreen, B. and H.C. Yee, "On Skew-Symmetric Splitting and Entropy Conservation Schemes for the Euler Equations," *8th European Conference on Numerical Mathematics and Advanced Applications (ENUMATH)*, (Springer), (LLNL-CONF-417535).

Sjogreen, B., et al., "Parallel Performance of ADPDIS3D: A High Order Multiblock Overlapping Grid Solver for Hypersonic Turbulence," *21st International Conference on Parallel CFD*, (DEStech Publications, Inc.), (LLNL-CONF-417430).

Springer, H.K., et al., "Satellite Collision Modeling with Physics-Based Hydrocodes: Debris Generation Predictions of the Iridium-Cosmos Collision Event and Other Impact Events," *11th Annual Advance Maui Optical and Space Surveillance Technologies Conference*.

Springer, H.K., et al., "Mesoscale Modeling of LX-17 Under Isentropic Compression," *14th International Detonation Symposium*, (LLNL-CONF-425427).

Vassilevski, P.S., "General Constrained Energy Minimization Interpolation Mappings for AMG," *10th Copper Mountain Conference on Iterative Methods*, (SIAM J. Sci. Comput., SIAM Publications), **32**, 1–13.

Vo, A., et al., "A Scalable and Distributed Dynamic Formal Verifier for MPI Programs," *SC10*, (IEEE and ACM), (LLNL-CONF-442475).

Walker, H., C. Woodward, and U.M. Yang, "An Accelerated Fixed-Point Iteration for Solution of Variably Saturated Flow," *28th International Conference on Water Resources (CMWR 2010)*.

Wang, P., O.A. Hurricane, M. Zika, and D.T. Woods, "3D Numerical Simulations of Radiation Hydrodynamics in a System of Multimaterials (U)," *Nuclear Explosives Code Development Conference (NECDC)*.

Wang, P., et al., "Simulation of Key 3D Aspects in Energy Balance (U)," *Nuclear Explosives Design Physics Conference (NEDPC)*.

Yee, H.C. and B. Sjogreen, "High Order Filter Methods for a Wide Range of Compressible Flow Speeds," *International Conference on Spectral and High Order Methods (ICOSAHOM)*, (Springer), (LLNL-CONF-417575).

Yee, H.C., B. Sjogreen, and A. Hadjadj, "Comparative Study of High Order Schemes for LES of Temporally Evolving Mixing Layers," *ASTRONUM-2010*.

Yoo, A., "Scalable Analysis of Massive Graphs on a Parallel Data Flow System," *43rd Hawaii International Conference on Systems Sciences (HICSS 2010)*, (IEEE Computer Society), 1492–1497.

COMPANY	TOPIC	LLNL CONTACT
Adaptive Computing Enterprises, Inc.	Moab workload manager	Don Lipari
Allinea Software	Scalable debugging infrastructure	Greg Lee
Appro	Scalable capacity clusters	Mark Seager
Aramco	Oil well reservoir visualization	Eric Brugger
Arista Networks	Low-latency Ethernet networks	Matt Leininger
Bull, Commissariat à l'Énergie Atomique, Hewlett-Packard, National Supercomputer Centre (Sweden), National University of Defense Technology (China), Swiss National Computer Centre	SLURM resource management software	Morris Jette
Cisco Systems	Application engine in router	Maya Gokhale
Cisco Systems, Dell, DataDirect Networks, Intel, LSI Industries, Mellanox, Qlogic, Red Hat, Sun Microsystems, Supermicro	Hyperion collaboration	Mark Seager
Cray	Programming environments	Bronis de Supinski
Cray	Scalable systems, open-source strategy	Mark Seager
DataDirect Networks	RAID 6 R&D for I/O systems	Keith Fitzgerald
Dell Computers	Scalable capacity clusters	Matt Leininger
Fusion-io	Solid-state memory	Maya Gokhale
IBM	High-performance storage system	Jerry Shoopman
IBM	Infosphere streams	Maya Gokhale
IBM	Scalable systems, multiple areas	Mark Seager
Imigene	DNA signatures for blood-borne diseases	Tom Slezak
Impulse Accelerated Technologies	C-to-FPGA compiler	Maya Gokhale
Intel Corporation	Solid-state memory	Maya Gokhale

INDUSTRIAL COLLABORATORS

APPENDIX C

COMPANY	TOPIC	LLNL CONTACT
Krell Institute	OpenSpeedShop performance analysis tool	Martin Schulz
LexisNexis Special Services, Inc.	Data analytic supercomputer applications	Steven Bradley
Microfluidic Systems	Pathogen signatures	Tom Slezak
Nallatech	Reconfigurable computing board	Maya Gokhale
Navistar Corporation, Michelin	Heavy vehicle aerodynamic drag reduction	Kambiz Salari
Northrup Grumman	DHS Domestic Nuclear Detection Office Red Team HSDN collaborative Web presence	John Brown
NVIDIA	3D-mesh compression	Jon Cohen
OpenFabrics Alliance, Mellanox, QLogic	OpenFabrics enterprise distribution	Ira Weiny
OpenMP Consortium	Shared-memory programming models	Bronis de Supinski
OpenWorks	Valgrind memory tool and threading tool development	John Gyllenhaal
Oracle	Lustre file system development and deployment	Chris Morrone and Marc Stearman
Palantir Technologies	DOE NA-22 Interdiction Technical Analysis Group NPAIR project	John Brown
Red Hat	Operating systems	Mark Grondona
Rogue Wave Software	TotalView parallel debugger scalability and enhanced memory tools	Dong Ahn and Scott Futral
Schlumberger	Seismic reverse time migration	Shawn Larsen
Siemens	Wind power	Bill Henshaw
SRI International	Acoustic and meteorological modeling	Shawn Larsen
WhamCloud	Lustre file system development and support	Kim Cupps and Chris Morrone
Xiotech	High performance I/O subsystems	Matt Leininger
Virident	Storage class memory	Maya Gokhale

RETURN TO TOC

Cover design by Amy Henke

About the Cover

The artist's conception on the cover features a circuit board, superimposed on binary code and several "grand challenge" equations, a selection of the expressions that relate fundamental quantities in the various computational sciences.

Production Team

LLNL Associate Director for Computation
Dona Crawford

Deputy Associate Directors
Ronald Schmucker (Principal Deputy), David L. Brown, John Grosh, and Michel McCoy

Scientific Editors
Tyler J. Alumbaugh and Susan Springer

Art Director
Amy E. Henke

Production Designer
Daniel Moore

Production Editor
Deanna Willis

Editorial Assistance
Caryn Meissner

Proofreader
Kris Fury

Photographers
Lee Baker and Nancy Rutter

Print Production
Charlie Arteago, Jr., and Monarch Print Copy and Design Solutions

LLNL-TR-470873